

Weak MSO: Automata and Expressiveness Modulo Bisimilarity*

Facundo Carreiro

ILLC, Universiteit van Amsterdam, The Netherlands

Alessandro Facchini

University of Warsaw, Poland

Yde Venema

ILLC, Universiteit van Amsterdam, The Netherlands

Fabio Zanasi

ENS Lyon, U. Lyon, CNRS, INRIA, UCBL, France

Last revision: January 23, 2014

Abstract

We prove that the bisimulation-invariant fragment of weak monadic second-order logic (WMSO) is equivalent to the fragment of the modal μ -calculus where the application of the least fixpoint operator $\mu p.\varphi$ is restricted to formulas φ that are continuous in p . Our proof is automata-theoretic in nature; in particular, we introduce a class of automata characterizing the expressive power of WMSO over tree models of arbitrary branching degree. The transition map of these automata is defined in terms of a logic FOE_1^∞ that is the extension of first-order logic with a generalized quantifier \exists^∞ , where $\exists^\infty x.\varphi$ means that there are infinitely many objects satisfying φ . An important part of our work consists of a model-theoretic analysis of FOE_1^∞ .

1 Introduction

Expressiveness modulo bisimilarity. This paper concerns the relative expressive power of some languages used for describing properties of pointed labelled transitions systems, or Kripke models. The interest in such expressiveness questions stems from applications where these structures model computational processes, and bisimilar pointed structures represent the *same* process. Seen from this perspective, properties of transition structures are relevant only if they are invariant under bisimilarity. This explains the importance of bisimulation invariance results of the form

$$M \equiv L / \simeq \text{ (over } K \text{)}$$

*Emails: fcarreiro@dc.uba.ar, facchini@mimuw.edu.pl, y.venema@uva.nl, fabio.zanasi@ens-lyon.fr.

stating that, if one restricts attention to a certain class K of transition structures, one language M is expressively complete with respect to the relevant (i.e., bisimulation-invariant) properties that can be formulated in another language L . In this setting, generally L is some rich yardstick formalism such as first-order or monadic second-order logic, and M is some modal-style fragment of L , usually displaying much better computational behavior than the full language L .

A seminal result in the theory of modal logic is van Benthem’s Characterization Theorem [20], stating that every bisimulation-invariant first-order formula $\alpha(x)$ is actually equivalent to (the standard translation of) a modal formula:

$$\text{ML} \equiv \text{FO}/\equiv \text{ (over the class of all LTSs).}$$

Over the years, a wealth of variants of the Characterization Theorem have been obtained. For instance, Rosen proved that van Benthem’s theorem is one of the few preservation results that transfers to the setting of finite models [17]; for a recent, rich source of van Benthem-style characterization results, see Dawar & Otto [2]. In this paper we are mainly interested in the work of Janin & Walukiewicz [9], who extended van Benthem’s result to the setting of fixpoint logics, by proving that the modal μ -calculus (μML) is the bisimulation-invariant fragment of monadic second-order logic (MSO):

$$\mu\text{ML} \equiv \text{MSO}/\equiv \text{ (over the class of all LTSs).}$$

Bisimulation invariance for WMSO. The yardstick logic that we consider in this paper is *weak* monadic second-order logic (WMSO), a variant of monadic second-order logic where the second-order quantifiers range over *finite* subsets of the transition structure rather than over arbitrary ones. Our target will be to identify the bisimulation-invariant fragment of this logic WMSO.

Before moving on, we should stress the role of the ambient class K in bisimulation-invariance results. Of particular importance in the setting of weak monadic second-order logic is the difference between structures of finite versus arbitrary branching degree. In the case of finitely branching models, it is not very hard to show that WMSO is a (proper) fragment of MSO, and it seems to be folklore that WMSO/\equiv corresponds to AFMC, the alternation-free fragment of the modal μ -calculus. For binary trees, this result was proved by Arnold & Niwiński in [1]. In the case of structures of arbitrary branching degree, however, WMSO and MSO have *incomparable* expressive power. The fact that, in particular, WMSO does not correspond to a fragment of MSO, is witnessed by the class of infinitely branching structures, which is clearly WMSO-definable, cannot be defined in MSO, since every MSO-definable class of trees contains a finitely branching tree.¹ For this reason, the relative expressive power of WMSO/\equiv and MSO/\equiv is not a priori clear. However, it is reasonable to think that WMSO/\equiv is strictly *weaker* than AFMC: the class of well-founded trees, which is definable in AFMC by the simple formula $\mu p.\Box p$, is not definable in WMSO.² Incidentally,

¹As remarked in [18], this follows from the automata characterization of MSO given in [23].

²This follows from the fact that WMSO can only define properties of trees that, from a topological point of view, are Borel, which is not the case of the class of trees defined by $\mu p.\Box p$ — see e.g. [18].

the question whether, conversely, there is a natural logic of which the bisimulation-invariant fragment corresponds to AFMC was answered positively by three of the present authors in [5], where they introduced another variant of MSO, called well-founded MSO (WFMSO), and proved that $\text{WFMSO}/\equiv \equiv \text{AFMC}$ (over the class of all LTSs).

The main result that we shall prove in this paper states that the bisimulation-invariant fragment of WMSO is equivalent to a certain, fragment $\mu_c\text{ML}$ of the modal μ -calculus.

$$\mu_c\text{ML} \equiv \text{WMSO}/\equiv \text{ (over the class of all LTSs).} \quad (1)$$

This fragment $\mu_c\text{ML}$, which is strictly weaker than the alternation-free fragment of μML , is characterized by a certain restriction on the application of fixpoint operators, which involves the notion of (*Scott*) *continuity*.

Continuity, an interesting property that features naturally in the semantics of many (fix-point) logics, in fact plays a key role throughout this paper. For its definition, we consider how the meaning $\llbracket \varphi \rrbracket^{\mathbb{T}} \subseteq T$ of a formula φ in some structure \mathbb{T} (with domain T) depends on the meaning of a fixed proposition letter or monadic predicate symbol p . This dependence can be formalized as a map $\varphi_p^{\mathbb{T}} : \wp(T) \rightarrow \wp(T)$, and if this map satisfies the condition

$$\varphi_p^{\mathbb{T}}(X) = \bigcup \left\{ \varphi_a^{\mathbb{T}}(X') \mid X' \text{ is a finite subset of } X \right\}, \quad (2)$$

we say that φ is *continuous in p* . The topological terminology stems from the observation that (2) expresses the continuity of the map $\varphi_p^{\mathbb{T}}$ with respect to the Scott topology on $\wp(T)$. If we look at concrete cases, this definition can be given a different reading: if φ is a formula of the modal μ -calculus, (2) means that φ holds at some state s of \mathbb{T} iff we can shrink the interpretation of the proposition letter p to some finite subset of the original interpretation, in such a way that φ holds at s in the modified version of \mathbb{T} .

A syntactic *characterization* of this property for the modal μ -calculus was obtained by Fontaine [6, 7], and the definition of our fragment $\mu_c\text{ML}$ uses this characterization as follows: whereas in the full language of μML the only syntactic condition on the formation of a formula $\mu p.\varphi$ is that φ is *positive* in p , for the fragment $\mu_c\text{ML}$ this condition is strengthened to the requirement that φ is (syntactically) *continuous* in p . More precisely, the fragment $\mu_c\text{ML}$ is defined as follows:

Definition 1.1 For each set Q of proposition letters, the fragment μMLC_Q of μML which is *continuous in Q* is given by the simultaneous induction

$$\varphi ::= q \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond \varphi \mid \mu p.\alpha$$

where $p \in P$, $q \in Q$, ψ is a Q -free μML -formula, and $\alpha \in \mu\text{MLC}_{Q \cup \{p\}}$. The formulas of the fragment $\mu_c\text{ML}$ are then given by the following induction:

$$\varphi ::= p \mid \neg \varphi \mid \varphi \vee \varphi \mid \Diamond \varphi \mid \mu p.\alpha$$

where $p \in P$, and $\alpha \in \mu\text{MLC}_p$.

In fact we will prove, analogous to the result by Janin & Walukiewicz, the following strong version of the characterization result (1), which provides an explicit translation, mapping any bisimulation-invariant formula φ in WMSO to an equivalent formula φ^\bullet in $\mu_c\text{ML}$.

Theorem 1.2 *There are effective translations $(-)^\bullet : \text{WMSO} \rightarrow \mu_c\text{ML}$ and $(-)_{\bullet} : \mu_c\text{ML} \rightarrow \text{WMSO}$ such that*

1. *A formula φ of WMSO is bisimulation invariant if and only if $\varphi \equiv \varphi^\bullet$, and*
2. *$\psi \equiv \psi_{\bullet}$ for every formula $\psi \in \mu_c\text{ML}$.*

To see how this theorem implies (1), observe that part (i) shows that $\text{WMSO}/\equiv \leq \mu_c\text{ML}$. Part (ii) states that $\mu_c\text{ML} \leq \text{WMSO}$, so combined with the fact that every formula in $\mu_c\text{ML} \subseteq \mu\text{ML}$ is bisimulation invariant, this gives the converse, $\mu_c\text{ML} \leq \text{WMSO}/\equiv$.

Automata for WMSO. As usual in this research area, our proof will be automata-theoretic in nature. More specifically, as the second main contribution of this paper, we introduce a new class of parity automata that exactly captures the expressive power of WMSO over the class of tree models of arbitrary branching degree.

Before we turn to a description of these automata, we first have a look at the automata, introduced by Walukiewicz [23], corresponding to MSO (over tree models). Fixing the set of proposition letters of our models as P , we think of $\wp(P)$ as an *alphabet* or set of *colors*. We can then define an MSO-automaton as a tuple $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$, where A is a finite set of states, a_I an initial state, and $\Omega : A \rightarrow \mathbb{N}$ is a parity function. The transition function Δ maps a pair $(a, c) \in A \times \wp(P)$ to a sentence in the first-order language (with equality) $\text{FOE}_1(A)$, of which the state space A provides the set of (monadic) predicates. For a more precise definition, let $\text{FOE}_1^+(A)$ denote the set of those sentences in $\text{FOE}_1(A)$ where all predicates in A occur only positively; we require that $\Delta : A \times \wp(P) \rightarrow \text{FOE}_1^+(A)$.

We shall refer to FOE_1 as the *one-step language* of MSO-automata, and denote the class of MSO-automata with $\text{Aut}(\text{FOE}_1)$. The automata that we consider in this article run on labelled transition systems and decide whether to accept or reject them. To take such decision we associate an acceptance game for an MSO-automaton \mathbb{A} and a transition system \mathbb{T} . A match of this game consists of two players, \exists and \forall , moving a token from one position to another. When such a match arrives at a so-called *basic* position, i.e., a position of the form $(a, t) \in A \times T$, the players consider the sentence $\Delta(a, c_t) \in \text{FOE}_1^+(A)$, where $c_t \in \wp(P)$ is the color of t (that is, the set of proposition letters true at t). At this position \exists has to turn the set $R[t]$ of successors of s into a *model* for the formula $\Delta(a, c_t)$ by coming up with an interpretation I of the monadic predicates $a \in A$ as subsets of $R[s]$, so that the resulting first-order structure $(R[s], I)$ makes the formula $\Delta(a, c_t)$ true.

Walukiewicz's key result linking MSO to $\text{Aut}(\text{FOE}_1)$ states that

$$\text{MSO} \equiv \text{Aut}(\text{FOE}_1) \text{ (over tree models),} \quad (3)$$

and the proof of this result proceeds by inductively showing that every formula φ in MSO can be effectively transformed into an equivalent automaton $\mathbb{A}_\varphi \in \text{Aut}(\text{FOE}_1)$. For the details of

this construction, a fairly intricate analysis of the one-step logic FOE_1 is required, crucially involving various normal forms of the sentences of $\text{FOE}_1(A)$.

In order to adapt this approach to the setting of WMSO, observe that by König's lemma, a subset of a tree \mathbb{T} is finite iff it is both a subset of a finitely branching subtree of \mathbb{T} and *noetherian*, that is, a subset of a subtree of \mathbb{T} that has no infinite branches. This suggests that we may change the definition of MSO-automata into one of WMSO-automata via two kinds of modifications, roughly speaking corresponding to a horizontal and a vertical 'dimension' of trees.

For the 'vertical modification' we may turn to the literature on weak automata [14]. The acceptance condition Ω of a parity automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ is *weak* if $\Omega(a) = \Omega(a')$ whenever the states a and a' belong to the same strongly connected component (SCC) of the automaton. To see that the notion of connected component is well-defined observe that for \mathbb{A} we can associate a directed graph on A such that $a, b \in A$ are connected iff b occurs in $\Delta(a, c)$ for some $c \in \wp(P)$. Let $\text{Aut}_w(\text{FOE}_1)$ denote the set of MSO-automata with a weak parity condition. It was proved in [24] (see also [5]) that

$$\text{WFMSO} \equiv \text{Aut}_w(\text{FOE}_1) \text{ (over the class of all trees),}$$

with WFMSO denoting the earlier mentioned variant of MSO where second-order quantification is restricted to noetherian subsets of trees. From this it easily follows that

$$\text{WMSO} \equiv \text{Aut}_w(\text{FOE}_1) \text{ (over the class of finitely branching trees),}$$

since the noetherian subsets of a finitely branching trees correspond to the finite ones. Over the class of all tree models, however, WMSO is *not* equivalent to $\text{Aut}_w(\text{FOE}_1)$, as is witnessed by the earlier mentioned class of well-founded trees, which can be defined in $\text{AFMC} \leq \text{WFMSO}$, but not in WMSO.

The hurdle to take, in order to find automata for WMSO on trees of *arbitrary* branching degree, concerns the horizontal dimension; the main problem lies in finding the right one-step language for WMSO-automata. An obvious candidate for this language would be weak monadic second-order logic itself, or more precisely, its variant WMSO_1 over the signature of monadic predicates (corresponding to the automata states). A very helpful observation, made by Väänänen [19], states that

$$\text{WMSO}_1 \equiv \text{FOE}_1^\infty,$$

where FOE_1^∞ is the extension of FOE_1 with the generalized quantifier \exists^∞ , where $\exists^\infty x.\varphi$ meaning that there are *infinitely* many objects satisfying φ . Taking the *full* language of WMSO_1 or FOE_1^∞ as our one-step language would give too much expressive power: since FOE_1^∞ extends FOE_1 , we would find that, over tree models, $\text{Aut}_w(\text{FOE}_1^\infty)$ extends $\text{Aut}_w(\text{FOE}_1)$, whereas we already saw that $\text{Aut}_w(\text{FOE}_1) \equiv \text{WFMSO}$ is incomparable to WMSO. It is here that we will crucially involve the notion of *continuity*. The automata corresponding to WMSO will be of the form $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$, where the transition map $\Delta : A \times \wp(P) \rightarrow \text{FOE}_1^{\infty+}(A)$ is subject to the following two constraints, for all $a, a' \in A$ belonging to the same strongly connected component of A :

(**weakness**) $\Omega(a) = \Omega(a')$, and

(**continuity**) if $\Omega(a)$ is odd (resp. even), then for each colour $c \in \wp(\mathbf{P})$, $\Delta(a, c)$ is continuous (resp. co-continuous) in a' ,

where co-continuity is a dual notion to continuity. The class of these automata is denoted by $Aut_{cw}(\text{FOE}_1^\infty)$. Consequently, for a proper definition of these automata we need a *syntactic* characterization of the $\text{FOE}_1^\infty(A)$ -sentences that are (co-)continuous in one (or more) monadic predicates of A .

For this purpose, we conduct a fairly detailed model-theoretic study of the logic FOE_1^∞ which we consider to be the third main contribution of our work. Similar to the results for FOE_1 , we provide normal forms for the sentences of $\text{FOE}_1^\infty(A)$, and syntactic characterizations of the fragments whose sentences are monotone (respectively continuous) in some monadic predicate $a \in A$.

To finish, we give constructions transforming WMSO-formulas to WMSO-automata and vice-versa, witnessing that

$$\text{WMSO} \equiv Aut_{cw}(\text{FOE}_1^\infty) \text{ (over tree models)}. \quad (4)$$

Proof of main result. To conclude our introduction we briefly sketch the proof of our main result, Theorem 1.2(1). Roughly speaking, we follow the bisimulation-invariance proof by Janin & Walukiewicz, which revolves around relating two distinct types of automata, which correspond, respectively, to the logics MSO and μML . More precisely, these two automaton types are given as $Aut(\text{FOE}_1)$ and $Aut(\text{FO}_1)$, where the one-step languages are first-order logic respectively with and without equality. What we will add to their proof is the insight from [21] that the required relation between $Aut(\text{FOE}_1)$ and $Aut(\text{FO}_1)$ already follows from results relating the respective one-step languages.

In our setting, we need to identify automata corresponding to the fragment $\mu_c\text{ML}$. For this purpose we introduce the class $Aut_{cw}(\text{FO}_1)$ consisting of those automata in $Aut(\text{FO}_1)$ that satisfy similar weakness and continuity conditions as the ones in $Aut_{cw}(\text{FOE}_1^\infty)$:

$$\mu_c\text{ML} \equiv Aut_{cw}(\text{FO}_1) \text{ (over the class of all LTSs)}. \quad (5)$$

As the key step in our proof then, we will provide a translation $(-)^{\bullet} : \text{FOE}_1^\infty \rightarrow \text{FO}_1$ which naturally induces a transformation $(-)^{\bullet} : Aut_{cw}(\text{FOE}_1^\infty) \rightarrow Aut_{cw}(\text{FO}_1)$. As a consequence of the nice model-theoretic properties of the translation at the one-step level, the automaton transformation satisfies, for all transition systems \mathbb{T} :

$$\mathbb{A}^{\bullet} \text{ accepts } \mathbb{T} \text{ iff } \mathbb{A} \text{ accepts } \mathbb{T}^\omega \quad (6)$$

where \mathbb{T}^ω is the ω -unravelling of \mathbb{T} . It easily follows from (6) that a WMSO-automaton \mathbb{A} is bisimulation invariant iff $\mathbb{A} \equiv \mathbb{A}^{\bullet}$, and so Theorem 1.2(1) follows by (4) and (5).

Overview of paper. In the next section we give a precise definition of the preliminaries required to understand this article. In Section 3 we define the one-step logics that will be used through the paper and give normal forms and syntactic characterizations of their

monotone and (co-)continuous fragments. In Section 4 we formally define WMSO-automata and show that from every WMSO-formula we can construct an equivalent WMSO-automaton. In Section 5 we prove the converse, that is, for every WMSO-automaton we can construct an equivalent WMSO-formula, this finishes the automata characterization of WMSO over tree models. Finally, in Section 6 we prove the main result of the paper, namely that the fragment $\mu_c\text{ML}$ is the bisimulation-invariant fragment of WMSO.

2 Preliminaries

2.1 Transition systems and trees

Throughout this article we fix a set P of elements that will be called *proposition letters* and denoted with small Latin letters p, q, \dots . We denote with C the set $\wp(P)$ of *labels* on P ; it will be convenient to think of C as an *alphabet*. Given a binary relation $R \subseteq X \times Y$, for any element $x \in X$, we indicate with $R[x]$ the set $\{y \in Y \mid (x, y) \in R\}$ while R^+ and R^* are defined respectively as the transitive closure of R and the reflexive and transitive closure of R . The set $\text{Ran}(R)$ is defined as $\bigcup_{x \in X} R[x]$.

A *C-labeled transition system* (LTS) is a tuple $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ where T is the universe or domain of \mathbb{T} , $\sigma : T \rightarrow \wp(P)$ is a marking, $R \subseteq T^2$ is the accessibility relation and $s_I \in T$ is a distinguished node. We use $|\mathbb{T}|$ to denote the domain of \mathbb{T} . Observe that the marking $\sigma : T \rightarrow \wp(P)$ can be seen as a valuation $\sigma^b : P \rightarrow \wp(T)$ given by $\sigma^b(p) = \{s \in T \mid p \in \sigma(s)\}$.

A *C-tree* is a LTS in which every node can be reached from s_I , and every node except s_I has a unique predecessor; the distinguished node s_I is called the *root* of \mathbb{T} . Each node $s \in T$ uniquely defines a subtree of \mathbb{T} with carrier $R^*[s]$ and root s . We denote this subtree by $\mathbb{T}.s$. We use the term *tree language* as a synonym of class of *C-trees*.

The tree unravelling of an LTS \mathbb{T} is given by $\mathbb{T}^e := \langle T_P, R_P, \sigma', s_I \rangle$ where T_P is the set of finite paths in \mathbb{T} stemming from s_I , $R_P(t, t')$ iff t' is an extension of t and the color of a path $t \in T_P$ is given by the color of its last node in T . The ω -unravelling \mathbb{T}^ω of \mathbb{T} is an unravelling which has ω -many copies of each node different from the root.

A *p-variant* of a transition system $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ is a $\wp(P \cup \{p\})$ -transition system $\langle T, R, \sigma', s_I \rangle$ such that $\sigma'(s) \setminus \{p\} = \sigma(s)$ for all $s \in T$. Given a set $S \subseteq T$, we let $\mathbb{T}[p \mapsto S]$ denote the *p-variant* where $p \in \sigma'(s)$ iff $s \in S$.

Let $\varphi \in \mathcal{L}$ be a formula of some logic \mathcal{L} , we use $\llbracket \varphi \rrbracket = \{\mathbb{T} \mid \mathbb{T} \models \varphi\}$ to denote the class of transition systems that make φ true. A class K of transition systems is said to be \mathcal{L} -*definable* if there is a formula $\varphi \in \mathcal{L}$ such that $\llbracket \varphi \rrbracket = K$. We use the notation $\varphi \equiv \psi$ to mean that $\llbracket \varphi \rrbracket = \llbracket \psi \rrbracket$ and given two logics $\mathcal{L}, \mathcal{L}'$ we use $\mathcal{L} \equiv \mathcal{L}'$ when the \mathcal{L} -definable and \mathcal{L}' -definable classes of models coincide.

Convention. Throughout this paper, we will only consider transition systems \mathbb{T} in which $R[s]$ is non-empty for every node $s \in T$. In particular this means that every tree we consider is *leafless*. All our results, however, can easily be lifted to the general case.

2.2 Games

We introduce some terminology and background on infinite games. All the games that we consider involve two players called *Éloise* (\exists) and *Abelard* (\forall). In some contexts we refer to a player Π to specify a generic player in $\{\exists, \forall\}$. Given a set A , by A^* and A^ω we denote respectively the set of words (finite sequences) and streams (or infinite words) over A .

A *board game* \mathcal{G} is a tuple $(G_\exists, G_\forall, E, \text{Win})$, where G_\exists and G_\forall are disjoint sets whose union $G = G_\exists \cup G_\forall$ is called the *board* of \mathcal{G} , $E \subseteq G \times G$ is a binary relation encoding the *admissible moves*, and $\text{Win} \subseteq G^\omega$ is a *winning condition*. An *initialized board game* $\mathcal{G}@u_I$ is a tuple $(G_\exists, G_\forall, u_I, E, \text{Win})$ where $u_I \in G$ is the *initial position* of the game. When Win is given by a parity function $\Omega : G \rightarrow \omega$ we say that \mathcal{G} is a parity game and sometimes simply write $\mathcal{G} = (G_\exists, G_\forall, E, \Omega)$.

Given a board game \mathcal{G} , a *match* of \mathcal{G} is simply a path through the graph (G, E) ; that is, a sequence $\pi = (u_i)_{i < \alpha}$ of elements of G , where α is either ω or a natural number, and $(u_i, u_{i+1}) \in E$ for all i with $i + 1 < \alpha$. A match of $\mathcal{G}@u_I$ is supposed to start at u_I . Given a finite match $\pi = (u_i)_{i < k}$ for some $k < \omega$, we call $\text{last}(\pi) := u_{k-1}$ the *last position* of the match; the player Π such that $\text{last}(\pi) \in G_\Pi$ is supposed to move at this position, and if $E[\text{last}(\pi)] = \emptyset$, we say that Π *got stuck* in π . A match π is called *total* if it is either finite, with one of the two players getting stuck, or infinite. Matches that are not total are called *partial*. Any total match π is *won* by one of the players: If π is finite, then it is won by the opponent of the player who gets stuck. Otherwise, if π is infinite, the winner is \exists if $\pi \in \text{Win}$, and \forall if $\pi \notin \text{Win}$.

Given a board game \mathcal{G} and a player Π , let PM_Π^G denote the set of partial matches of \mathcal{G} whose last position belongs to player Π . A *strategy* for Π is a function $f : \text{PM}_\Pi^G \rightarrow G$. A match $\pi = (u_i)_{i < \alpha}$ of \mathcal{G} is *f-guided* if for each $i < \alpha$ such that $u_i \in G_\Pi$ we have that $u_{i+1} = f(u_0, \dots, u_i)$. Let $u \in G$ and a f be a strategy for Π . We say that f is a *surviving strategy* for Π in $\mathcal{G}@u$ if

- (i) For each f -guided partial match π of $\mathcal{G}@u$, if $\text{last}(\pi)$ is in G_Π then $f(\pi)$ is legitimate, that is, $(\text{last}(\pi), f(\pi)) \in E$.

We say that f is a *winning strategy* for Π in $\mathcal{G}@u$ if, additionally,

- (ii) Π wins each f -guided total match of $\mathcal{G}@u$.

If Π has a winning strategy for $\mathcal{G}@u$ then u is called a *winning position* for Π in \mathcal{G} . The set of positions of \mathcal{G} that are winning for Π is denoted by $\text{Win}_\Pi(\mathcal{G})$. A strategy f is called *positional* if $f(\pi) = f(\pi')$ for each $\pi, \pi' \in \text{Dom}(f)$ with $\text{last}(\pi) = \text{last}(\pi')$. A board game \mathcal{G} with board G is *determined* if $G = \text{Win}_\exists(\mathcal{G}) \cup \text{Win}_\forall(\mathcal{G})$, that is, each $u \in G$ is a winning position for one of the two players.

Fact 2.1 (Positional Determinacy of Parity Games [4, 13]) *For each parity game \mathcal{G} , there are positional strategies f_\exists and f_\forall respectively for player \exists and \forall , such that for every position $u \in G$ there is a player Π such that f_Π is a winning strategy for Π in $\mathcal{G}@u$.*

From now on, we always assume that each strategy we work with in parity games is positional. Moreover, we think of a positional strategy f_Π for player Π as a function $f_\Pi : G_\Pi \rightarrow G$.

2.3 Parity automata

We recall the definition of a parity automaton, adapted to our setting. Since we will be comparing parity automata defined in terms of various one-step languages, it makes sense to make the following abstraction.

Definition 2.2 Given a set A , we define an A -structure to be a pair (D, V) consisting of a domain D and a valuation $V : A \rightarrow \wp D$. Depending on context, elements of A will be called *monadic predicates* or *propositional variables*.

A *one-step language* is a map \mathcal{L}_1 assigning to each set A a set $\mathcal{L}_1(A)$ of objects called *one-step formulas* over A . We require that $\mathcal{L}_1(\bigcap_i A_i) = \bigcap_i \mathcal{L}_1(A_i)$, so that for each $\varphi \in \mathcal{L}_1(A)$ there is a smallest $A_\varphi \subseteq A$ such that $\varphi \in \mathcal{L}_1(A_\varphi)$; this A_φ is the set of propositional variables that *occur* in φ .

We assume that such a one-step language \mathcal{L}_1 comes with a *truth* relation: given an A -structure (D, V) , a formula $\varphi \in \mathcal{L}_1$ is either *true* or *false* in (D, V) , denoted by, respectively, $(D, V) \models \varphi$ and $(D, V) \not\models \varphi$. We also assume that \mathcal{L}_1 has a *positive fragment* \mathcal{L}_1^+ characterizing monotonicity in the sense that a formula $\varphi \in \mathcal{L}_1(A)$ is (semantically) monotone iff it is equivalent to a formula $\varphi' \in \mathcal{L}_1^+(A)$.

The one-step languages \mathcal{L}_1 featuring in this paper all are induced by well-known logics. Examples include first-order logic (with and without equality), first-order logic extended with the infinity quantifier, and fragments of these languages.

Definition 2.3 Let \mathcal{L}_1 be some one-step language. A *parity automaton* based on \mathcal{L}_1 and alphabet C is a tuple $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ such that A is a finite set of states of the automaton, $a_I \in A$ is the initial state, $\Delta : A \times C \rightarrow \mathcal{L}_1^+(A)$ is the transition map, and $\Omega : A \rightarrow \mathbb{N}$ is the parity map. The collection of such automata will be denoted by $\text{Aut}(\mathcal{L}_1)$.

Acceptance and rejection of a transition system by an automaton is defined in terms of the following parity game.

Definition 2.4 Given $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $\text{Aut}(\mathcal{L}_1)$ and a transition system $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$, the *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{T})$ of \mathbb{A} on \mathbb{T} is the parity game defined according to the rules of the following table.

Position	Player	Admissible moves	Parity
$(a, s) \in A \times T$	\exists	$\{V : A \rightarrow \wp(R[s]) \mid (R[s], V) \models \Delta(a, \sigma(s))\}$	$\Omega(a)$
$V : A \rightarrow \wp(T)$	\forall	$\{(b, t) \mid t \in V(b)\}$	$\max(\Omega[A])$

A transition system \mathbb{T} is *accepted* by \mathbb{A} if \exists has a winning strategy in $\mathcal{A}(\mathbb{A}, \mathbb{T})@ (a_I, s_I)$, and *rejected* if (a_I, s_I) is a winning position for \forall .

Many properties of parity automata are determined at the one-step level. An important example concerns the notion of complementation.

Definition 2.5 Two one-step formulas φ and ψ are each other's *Boolean dual* if for every structure (D, V) we have

$$(D, V) \models \varphi \text{ iff } (D, V^c) \not\models \psi,$$

where V^c is the valuation given by $V^c(a) := D \setminus V(a)$, for all a . A one-step language \mathcal{L}_1 is *closed under Boolean duals* if for every set A , each formula $\varphi \in \mathcal{L}_1(A)$ has a Boolean dual $\varphi^\delta \in \mathcal{L}_1(A)$.

Following ideas by [10], we can use Boolean duals, together with a *role switch* between \forall and \exists , in order to define a negation or complementation operation on automata.

Definition 2.6 Assume that, for some one-step language \mathcal{L}_1 , the map $(-)^\delta$ provides, for each set A , a Boolean dual $\varphi^\delta \in \mathcal{L}_1(A)$ for each $\varphi \in \mathcal{L}_1(A)$. Given $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $\text{Aut}(\mathcal{L}_1)$ we define its *complement* $\overline{\mathbb{A}}$ as the automaton $\langle A, \Delta^\delta, \Omega^\delta, a_I \rangle$, where $\Delta^\delta(a, c) := (\Delta(a, c))^\delta$, and $\Omega^\delta(a) := 1 + \Omega(a)$, for all $a \in A$ and $c \in C$.

Proposition 2.7 Let \mathcal{L}_1 and $(-)^\delta$ be as in the previous definition. For each automaton $\mathbb{A} \in \text{Aut}(\mathcal{L}_1)$ and each transition structure \mathbb{T} we have that

$$\overline{\mathbb{A}} \text{ accepts } \mathbb{T} \text{ iff } \mathbb{A} \text{ rejects } \mathbb{T}.$$

The proof of Proposition 2.7 is based on the fact that the power of \exists in $\mathcal{A}(\overline{\mathbb{A}}, \mathbb{T})$ is the same as that of \forall in $\mathcal{A}(\mathbb{A}, \mathbb{T})$.

As an immediate consequence of this proposition, one may show that if the one-step language \mathcal{L}_1 is closed under Boolean duals, then the class $\text{Aut}(\mathcal{L}_1)$ is closed under taking complementation. Further on we will use Proposition 2.7 to show that the same may apply to some subsets of $\text{Aut}(\mathcal{L}_1)$.

2.4 Weak monadic second-order logic

Definition 2.8 The *weak monadic second-order logic* on a set of predicates P is given by

$$\varphi ::= \Downarrow p \mid p \sqsubseteq q \mid R(p, q) \mid \neg \varphi \mid \varphi \vee \psi \mid \exists p. \varphi$$

where $p, q \in P$. We denote this logic by $\text{WMSO}(P)$ and omit P when the set of proposition letters is clear from context. We adopt the standard convention that no letter is both free and bound in φ .

Definition 2.9 Let $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ be a LTS, the semantics of WMSO is defined as follows

$$\begin{aligned} \mathbb{T} \models \Downarrow p & \text{ iff } \sigma^b(p) = \{s_I\} \\ \mathbb{T} \models p \sqsubseteq q & \text{ iff } \sigma^b(p) \subseteq \sigma^b(q) \\ \mathbb{T} \models R(p, q) & \text{ iff for every } s \in \sigma^b(p) \text{ there exists } t \in \sigma^b(q) \text{ such that } sRt \\ \mathbb{T} \models \neg \varphi & \text{ iff } \mathbb{T} \not\models \varphi \\ \mathbb{T} \models \varphi \vee \psi & \text{ iff } \mathbb{T} \models \varphi \text{ or } \mathbb{T} \models \psi \\ \mathbb{T} \models \exists p. \varphi & \text{ iff there is a finite set } X \subseteq_\omega T \text{ such that } \mathbb{T}[p \mapsto X] \models \varphi \end{aligned}$$

Remark 2.10 The reader may have expected a more standard two-sorted language for second-order logic, for example given by

$$\varphi ::= p(x) \mid R(x, y) \mid x \approx y \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists x.\varphi \mid \exists p.\varphi$$

where $p \in \mathbf{P}$, $x, y \in \mathbf{iVar}$ (individual variables), and \approx is the symbol for equality. Both definitions can be proved to be equivalent, however, we choose to keep Definition 2.9 as it will be better suited to work with in the context of automata. \triangleleft

2.5 First-order logic with generalized quantifiers

In this subsection we introduce an extension of first-order logic with so called generalized quantifiers. Our interest in this extension stems from the fact that it will allow us to define a variant of first-order logic that is expressively equivalent to weak monadic second order (see Section 3) and has nice technical features such as a normal form theorem.

Mostowski [12] defined (unary) generalized quantifiers as follows: a (unary) generalized quantifier \mathcal{Q} is a collection of pairs (J, X) with $X \subseteq J$, and satisfying the following condition

$$\text{If } ((J, X) \in \mathcal{Q}, |X| = |Y| \wedge |J \setminus X| = |J \setminus Y|) \text{ then } (J, Y) \in \mathcal{Q}.$$

The semantics of \mathcal{Q} is defined by the following condition

$$\mathbb{T} \models \mathcal{Q}x.\varphi(x, \vec{s}) \text{ iff } (T, \{t \mid \mathbb{T} \models \varphi(t, \vec{s})\}) \in \mathcal{Q}.$$

where we use $\vec{s} := [s_1, \dots, s_n]$ to denote sequences or vectors of elements.

For the rest of this article we will focus on the generalized quantifier \exists^∞ expressing that there exist infinitely many elements satisfying a certain condition. Formally, it is defined as

$$\exists^\infty := \{(J, X) \mid |X| \geq \aleph_0\}.$$

The dual of \exists^∞ is $\forall^\infty = \{(J, X) \mid |J \setminus X| < \aleph_0\}$. It is worth observing what is the intended meaning of this quantifier: $\forall^\infty x.\varphi$ expresses that there are *at most finitely many* elements *falsifying* the formula φ . The extension of first-order logic with equality (FOE) or without equality (FO) obtained by adding \exists^∞ to the corresponding first-order language is denoted respectively by FO^∞ and FOE^∞ .

2.6 The Modal μ -Calculus.

The language of the modal μ -calculus (μML) on \mathbf{P} is given by the following grammar:

$$\varphi ::= q \mid \neg q \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Diamond\varphi \mid \Box\varphi \mid \mu p.\varphi \mid \nu p.\varphi$$

where $p, q \in \mathbf{P}$ and p is positive in φ (i.e., p is not negated). We use the standard convention that no variable is both free and bound in a formula and that every bound variable is fresh. Let p be a bound variable occurring in some formula $\varphi \in \mu\text{ML}$, we use δ_p to denote the binding definition of p , that is, the formula such that either $\mu p.\delta_p$ or $\nu p.\delta_p$ are subformulas of φ .

The semantics of this language is completely standard. Let $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ be a transition system and $\varphi \in \mu\text{ML}$. We inductively define the *meaning* $\llbracket \varphi \rrbracket^{\mathbb{T}}$ which includes the following clauses for the least (μ) and greatest (ν) fixpoint operators:

$$\begin{aligned}\llbracket \mu x. \psi \rrbracket^{\mathbb{T}} &:= \bigcap \{ S \subseteq T \mid S \supseteq \llbracket \psi \rrbracket^{\mathbb{T}[x \mapsto S]} \} \\ \llbracket \nu x. \psi \rrbracket^{\mathbb{T}} &:= \bigcup \{ S \subseteq T \mid S \subseteq \llbracket \psi \rrbracket^{\mathbb{T}[x \mapsto S]} \}\end{aligned}$$

We say that φ is *true* in \mathbb{T} (notation $\mathbb{T} \models \varphi$) iff $s_I \in \llbracket \varphi \rrbracket^{\mathbb{T}}$.

We will now describe the semantics defined above in game-theoretic terms. That is, we will define the evaluation game $\mathcal{E}(\varphi, \mathbb{T})$ associated with a formula $\varphi \in \mu\text{ML}$ and a transition system \mathbb{T} . This game is played by two players (\exists and \forall) moving through positions (ξ, s) where ξ is a subformula of φ and $s \in T$. In an arbitrary position (ξ, s) it is useful to think of \exists

Position	Player	Admissible moves
$(\psi_1 \vee \psi_2, s)$	\exists	$\{(\psi_1, s), (\psi_2, s)\}$
$(\psi_1 \wedge \psi_2, s)$	\forall	$\{(\psi_1, s), (\psi_2, s)\}$
$(\Diamond \varphi, s)$	\exists	$\{(\varphi, t) \mid t \in R[s]\}$
$(\Box \varphi, s)$	\forall	$\{(\varphi, t) \mid t \in R[s]\}$
$(\mu p. \varphi, s)$	—	$\{(\varphi, s)\}$
$(\nu p. \varphi, s)$	—	$\{(\varphi, s)\}$
(p, s) with p bound in φ	—	$\{(\delta_p, s)\}$
$(\neg q, s) \in \mathbf{P} \times T$, q free in φ and $q \notin \sigma(s)$	\forall	\emptyset
$(\neg q, s) \in \mathbf{P} \times T$, q free in φ and $q \in \sigma(s)$	\exists	\emptyset
$(q, s) \in \mathbf{P} \times T$, q free in φ and $q \in \sigma(s)$	\forall	\emptyset
$(q, s) \in \mathbf{P} \times T$, q free in φ and $q \notin \sigma(s)$	\exists	\emptyset

Table 1: Evaluation game for the modal μ -calculus

trying to show that ξ is true at s , and of \forall of trying to convince her that ξ is false at s . The rules of the evaluation game are given in Table 1. Every finite match of this game is lost by the player that got stuck. To give a winning condition for an infinite match let p be, of the bound variables of φ that get unravelled infinitely often, the one that is the highest in the syntactic tree of φ . The winner of the match is \forall if p is a μ -variable and \exists if p is a ν -variable. We say that φ is true in \mathbb{T} iff \exists has a winning strategy in $\mathcal{E}(\varphi, \mathbb{T})$.

Formulas of the modal μ -calculus are classified according to their *alternation depth*, which roughly is given as the maximal length of a chain of nested alternating least and greatest fixpoint operators [15]. The *alternation-free fragment* of the modal μ -calculus (AFMC) is the collection of μML -formulas without nesting of least and greatest fixpoint operators.

Definition 2.11 Let φ be a formula of the modal μ -calculus. We say that $\varphi \in \text{AFMC}$ iff for all subformulas $\mu p. \psi_1$ and $\nu q. \psi_2$ we have that p is not free in ψ_2 and q is not free in ψ_1 .

It is not difficult to see that, over arbitrary transition systems, this fragment is less expressive than the whole μML . That is, there is a μML -formula φ such that $\llbracket \varphi \rrbracket$ is not AFMC-definable [16].

In order to properly define the fragment $\mu_c\text{ML} \subseteq \text{AFMC}$ which is of critical importance in this article, we are particularly interested in the *continuous* fragment of the modal μ -calculus. As observed in Section 1, the abstract notion of continuity can be given a concrete interpretation in the context of μ -calculus.

Definition 2.12 Let $\varphi \in \mu\text{ML}$, and q be a propositional variable. We say that φ is *continuous in q* iff for every transition system \mathbb{T} there exists some finite $S \subseteq_\omega \sigma^b(q)$ such that

$$\mathbb{T} \models \varphi \quad \text{iff} \quad \mathbb{T}[q \mapsto S] \models \varphi.$$

We can give a syntactic characterization of the fragment of μML that captures this property. Given a set Q of propositional variables, we define the fragment of μML *continuous in Q* , denoted by μMLC_Q , by induction in the following way

$$\varphi ::= q \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond \varphi \mid \mu p. \alpha$$

where $q \in Q$, $p \in P \setminus Q$, $\alpha \in \mu\text{MLC}_{Q \cup \{p\}}$, and ψ is a Q -free μML -formula.

Proposition 2.13 ([6, 7]) *A μML -formula is continuous in q iff it is equivalent to a formula in the fragment μMLC_q .*

Finally, we define $\mu_c\text{ML}$ to be the fragment of μML where the use of the least fixed point operator is restricted to the continuous fragment. Formally, it is defined as follows.

Definition 2.14 Formulas of the fragment $\mu_c\text{ML}$ are given by the following induction:

$$\varphi ::= q \mid \neg \varphi \mid \varphi \vee \varphi \mid \Diamond \varphi \mid \mu p. \alpha$$

where $p, q \in P$, and $\alpha \in \mu\text{MLC}_p$.

Proposition 2.15 *Let $\varphi \in \mu_c\text{ML}$, the following hold*

- (1) φ is an AFMC-formula,
- (2) Every μ -variable in φ is existential (i.e., is only in the scope of diamonds), and dually every ν -variable in φ is universal (i.e., is only in the scope of boxes).

Proof. Both points are proved by an easy induction on the complexity of a formula. For the first one, it is enough to notice that if $\varphi \in \mu\text{MLC}_q \cap \text{AFMC}$, then $\mu q. \varphi \in \text{AFMC}$ by definition of μMLC_q . QED

As an immediate consequence of Proposition 2.13 we have the following:

Corollary 2.16 *For every $\mu_c\text{ML}$ -formula $\mu p. \varphi$, φ is continuous in p .*

2.7 Bisimulation

Bisimulation is a notion of behavioral equivalence between processes. For the case of transition systems, it is formally defined as follows.

Definition 2.17 Let $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ and $\mathbb{T}' = \langle T', R', \sigma', s'_I \rangle$ be transition systems. A *bisimulation* is a relation $Z \subseteq T \times T'$ such that for all $(t, t') \in Z$ the following holds:

- (**atom**) $p \in \sigma(t)$ iff $p \in \sigma'(t')$ for all $p \in \mathbf{P}$;
- (**forth**) for all $s \in R[t]$ there is $s' \in R'[t']$ such that $(s, s') \in Z$;
- (**back**) for all $s' \in R'[t']$ there is $s \in R[t]$ such that $(s, s') \in Z$.

Two pointed transition systems \mathbb{T} and \mathbb{T}' are *bisimilar* (denoted $\mathbb{T} \Leftrightarrow \mathbb{T}'$) if there is a bisimulation $Z \subseteq T \times T'$ containing (s_I, s'_I) .

The following fact about tree unravellings will allow us to provide a proof of Theorem 1.2 by just proving it for tree languages.

Fact 2.18 \mathbb{T} and \mathbb{T}^e are bisimilar, for every transition system \mathbb{T} .

A class of transition systems K is *bisimulation closed* if $\mathbb{T} \Leftrightarrow \mathbb{T}'$ implies that $\mathbb{T} \in K$ iff $\mathbb{T}' \in K$, for all \mathbb{T} and \mathbb{T}' . A formula $\varphi \in \mathcal{L}$ is *bisimulation-invariant* if $\mathbb{T} \Leftrightarrow \mathbb{T}'$ implies that $\mathbb{T} \models \varphi$ iff $\mathbb{T}' \models \varphi$, for all \mathbb{T} and \mathbb{T}' .

Fact 2.19 Each μML -definable class of transition systems is bisimulation closed.

3 One-step logics, normal forms and continuity

Given a set of names A we use $\mathcal{L}_1(A)$ to denote an arbitrary one-step language based on A . One-step formulas are interpreted over unary (or one-step) models $\mathbf{D} = (D, V : A \rightarrow \wp D)$. We use $\mathfrak{M}_1(A)$ to denote the class of all one-step models based on the set of names A .

In this section we define the one-steps logics that we use in the rest of the article, namely: one-step weak monadic second-order logic (WMSO_1), one-step first-order logic with and without equality (FOE_1 , FO_1) and one-step first-order logic with the \exists^∞ quantifier (FOE_1^∞). The main theorems of this section prove normal forms for these logics and give syntactical characterizations of the monotonic and continuous fragments that we use in later sections.

Definition 3.1 Let iVar be a set of (individual) variables. Given a set of names A we define the set $\text{WMSO}_1(A)$ of one-step weak monadic second-order formulas as the *sentences* given by the following grammar

$$\varphi ::= a(x) \mid x \approx y \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists x.\varphi \mid \exists a.\varphi$$

where $x, y \in \text{iVar}$, $a \in A$.

Definition 3.2 The set $\text{FOE}_1(A)$ of one-step first-order sentences (with equality) is given by the sentences formed by

$$\varphi ::= a(x) \mid x \approx y \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists x.\varphi$$

where $x, y \in \text{iVar}$, $a \in A$. The one-step logic $\text{FO}_1(A)$ is as $\text{FOE}_1(A)$ but without equality. The set $\text{FOE}_1^\infty(A)$ of one-step first-order sentences with generalized quantifier \exists^∞ (with equality) is defined analogously by just adding the clauses $\exists^\infty x.\varphi$ and $\forall^\infty x.\varphi$.

Without loss of generality, from now on we always assume that every bounded variable occurring in a sentence is bounded by an unique quantifier (generalized or not). Recall that given a one-step logic \mathcal{L}_1 we write $\mathcal{L}_1^+(A)$ to denote the fragment where every predicate $a \in A$ occurs only positively. The following observation will allow us to work with the (one-sorted) language FOE_1 instead of the (two-sorted) language WMSO_1 .

Fact 3.3 ([19]) $\text{WMSO}_1(A) \equiv \text{FOE}_1^\infty(A)$.

In the following subsections we provide a detailed model theoretic analysis of the one-step logics that we use in this article, specifically, we give

- Normal forms for arbitrary formulas of FO_1 , FOE_1 and FOE_1^∞ .
- Strong forms of syntactic characterizations for the monotone and continuous fragments of several of the mentioned logics. Namely, for $\mathcal{L}_1 \in \{\text{FO}_1, \text{FOE}_1, \text{FOE}_1^\infty\}$ we provide

- (a) A fragment $\mathcal{L}_1 M_a$ and a translation $(-)^\circ : \mathcal{L}_1(A) \rightarrow \mathcal{L}_1 M_a(A)$ such that for every $\varphi \in \mathcal{L}_1$ we have $\varphi \equiv \varphi^\circ$ iff φ is monotone in $a \in A$,

for $\mathcal{L}_1 \in \{\text{FO}_1, \text{FOE}_1^\infty\}$ we provide

- (b) A fragment $\mathcal{L}_1 C_a$ and a translation $(-)^\Delta : \mathcal{L}_1(A) \rightarrow \mathcal{L}_1 C_a(A)$ such that for every $\varphi \in \mathcal{L}_1$ we have $\varphi \equiv \varphi^\Delta$ iff φ is continuous in $a \in A$.

Moreover, we show that the latter translation also restricts to the fragment \mathcal{L}_1^+ , i.e.,

- (c) The restriction $(-)^\Delta_+ : \mathcal{L}_1^+(A) \rightarrow \mathcal{L}_1^+ C_a(A)$ of $(-)^\Delta$ is such that for every $\varphi \in \mathcal{L}_1^+$ we have $\varphi \equiv \varphi^\Delta_+$ iff φ is continuous in $a \in A$.

- Syntactic characterizations of the co-continuous fragments of FO_1 and FOE_1^∞ .
- Normal forms for the monotone and continuous fragments.

3.1 Normal forms

Given a set of names A and $S \subseteq A$, we introduce the notation

$$\tau_S(x) := \bigwedge_{a \in S} a(x) \wedge \bigwedge_{a \in A \setminus S} \neg a(x).$$

The formula $\tau_S(x)$ is called an *A-type*, we usually blur the distinction between $\tau_S(x)$ and S and call S an *A-type* as well. A *positive A-type* is defined as $\tau_S^+(x) := \bigwedge_{a \in S} a(x)$. We use the convention that, if S is the empty set, then $\tau_S^+(x)$ is \top and we call it an *empty positive A-type*. Given a one-step model \mathbf{D} we use $|S|_{\mathbf{D}}$ to denote the number of elements that realize the A-type τ_S in \mathbf{D} . Formally, it is defined as $|S|_{\mathbf{D}} := |\{d \in |\mathbf{D}| : \mathbf{D} \models \tau_S(d)\}|$.

A *partial isomorphism* between two one-step models $\mathbf{D} = (D, V)$ and $\mathbf{D}' = (D', V')$ is a *partial* function $f : D \rightarrow D'$ which is injective and satisfies $d \in V(a) \Leftrightarrow f(d) \in V'(a)$ for all $a \in A$ and $d \in \text{Dom}(f)$. Given two sequences $[d_1, \dots, d_k] \in D^k$ and $[d'_1, \dots, d'_k] \in D'^k$ we use $f : [d_1, \dots, d_k] \mapsto [d'_1, \dots, d'_k]$ to denote the partial function $f : D \rightarrow D'$ defined as $f(d_i) := d'_i$. If there exist d_i, d_j such that $d_i = d_j$ but $d'_i \neq d'_j$ then the result is undefined.

Definition 3.4 The quantifier rank $qr(\varphi)$ of $\varphi \in \text{FOE}_1^\infty$ (hence also for FO_1 and FOE_1) is defined as follows

- If φ is atomic $qr(\varphi) = 0$,
- If $\varphi = \neg\psi$ then $qr(\varphi) = qr(\psi)$,
- If $\varphi = \psi_1 \wedge \psi_2$ or $\varphi = \psi_1 \vee \psi_2$ then $qr(\varphi) = \max\{qr(\psi_1), qr(\psi_2)\}$,
- If $\varphi = Qx.\psi$ for $Q \in \{\exists, \forall, \exists^\infty, \forall^\infty\}$ then $qr(\varphi) = 1 + qr(\psi)$.

Given a one-step logic \mathcal{L}_1 we write $\mathbf{D} \equiv_k^{\mathcal{L}} \mathbf{D}'$ to indicate that the one-step models \mathbf{D} and \mathbf{D}' satisfy exactly the same formulas $\varphi \in \mathcal{L}_1$ with $qr(\varphi) \leq k$. The logic \mathcal{L} will be omitted when it is clear from context.

3.1.1 Normal form for FO_1

We start by stating a normal form for one-step first-order logic without equality. A formula in *basic form* gives a complete description of the types that are satisfied in a one-step model.

Definition 3.5 A formula $\varphi \in \text{FO}_1(A)$ is in *basic form* if $\varphi = \bigvee \nabla_{\text{FO}}(\Sigma)$ where each disjunct is of the form

$$\nabla_{\text{FO}}(\Sigma) = \bigwedge_{S \in \Sigma} \exists x. \tau_S(x) \wedge \forall x. \bigvee_{S \in \Sigma} \tau_S(x)$$

for some set of types $\Sigma \subseteq \wp A$.

It is easy to prove, using Ehrenfeucht-Fraïssé games, that every formula of first-order logic without equality over a unary signature (i.e., FO_1) is equivalent to a formula in basic form. Proof sketches can be found in [8, Lemma 16.23] and [21, Proposition 4.14]. We omit a full proof because it is very similar to the following cases.

Fact 3.6 Every formula of $\text{FO}_1(A)$ is equivalent to a formula in basic form.

3.1.2 Normal form for FOE₁

When considering a normal form for FOE₁, the fact that we can ‘count types’ using equality yields a more involved basic form.

Definition 3.7 A formula $\varphi \in \text{FOE}_1(A)$ is in *basic form* if $\varphi = \bigvee \nabla_{\text{FOE}}(\vec{T}, \Pi)$ where each disjunct is of the form

$$\nabla_{\text{FOE}}(\vec{T}, \Pi) = \exists \vec{x}. (\text{diff}(\vec{x}) \wedge \bigwedge_i \tau_{T_i}(x_i) \wedge \forall z. (\text{diff}(\vec{x}, z) \rightarrow \bigvee_{S \in \Pi} \tau_S(z)))$$

for some set of types $\Pi \subseteq \wp A$ and where each $T_i \subseteq A$. The predicate $\text{diff}(\vec{y})$, which states that the given elements are different, is defined as $\text{diff}(y_1, \dots, y_n) := \bigwedge_{1 \leq m < m' \leq n} (y_m \not\approx y_{m'})$.

We prove that every formula of monadic first-order logic with equality (i.e., FOE₁) is equivalent to a formula in basic form. This result seems to be folklore, however, we give a detailed proof because some of its ingredients are used in the following section, when giving a normal form for FOE₁[∞]. We start by defining the following relation between one-step models.

Definition 3.8 Let \mathbf{D} and \mathbf{D}' be one-step models. For every $k \in \mathbb{N}$ we define

$$\mathbf{D} \sim_k^\perp \mathbf{D}' \iff \forall S \subseteq A \left(|S|_{\mathbf{D}} = |S|_{\mathbf{D}'} < k \right. \\ \left. \text{or } |S|_{\mathbf{D}}, |S|_{\mathbf{D}'} \geq k \right)$$

Intuitively, two models are related by \sim_k^\perp when their type information coincides ‘modulo k ’. Later we will prove that this is the same as saying that they cannot be distinguished by a formula of FOE₁ with quantifier rank lower or equal to k . For the moment, we prove the following properties of \sim_k^\perp .

Proposition 3.9 *The following hold*

- (i) \sim_k^\perp is an equivalence relation,
- (ii) \sim_k^\perp has finite index,
- (iii) Every $E \in \mathfrak{M}_1 / \sim_k^\perp$ is characterized by a formula $\varphi_E^\perp \in \text{FOE}_1(A)$ with $qr(\varphi_E^\perp) = k$.

Proof. We only prove the last point. Let $E \in \mathfrak{M}_1 / \sim_k^\perp$ and let $\mathbf{D} \in E$ be a representative. Call $S_1, \dots, S_n \subseteq A$ to the types such that $|S_i|_{\mathbf{D}} = n_i < k$ and $S'_1, \dots, S'_m \subseteq A$ to those satisfying $|S'_i|_{\mathbf{D}} \geq k$. Now define

$$\varphi_E^\perp := \bigwedge_{i \leq n} (\exists x_1, \dots, x_{n_i}. \text{diff}(x_1, \dots, x_{n_i}) \wedge \bigwedge_{j \leq n_i} \tau_{S_i}(x_j) \wedge \forall z. \text{diff}(x_1, \dots, x_{n_i}, z) \rightarrow \neg \tau_{S_i}(z)) \wedge \\ \bigwedge_{i \leq m} (\exists x_1, \dots, x_k. \text{diff}(x_1, \dots, x_k) \wedge \bigwedge_{j \leq k} \tau_{S'_i}(x_j))$$

It is easy to see that $qr(\varphi_E^\perp) = k$ and that $\mathbf{D}' \models \varphi_E^\perp$ iff $\mathbf{D}' \in E$. Observe that φ_E^\perp gives a specification of E “type by type”. QED

In the following definition we recall a (standard) notion of Ehrenfeucht-Fraïssé game for FOE_1 which will be used to establish the connection between \sim_k^\perp and \equiv_k^{FOE} .

Definition 3.10 Let $\mathbf{D}_0 = (D_0, V_0)$ and $\mathbf{D}_1 = (D_1, V_1)$ be one-step models. We define the game $\text{EF}_k^\perp(\mathbf{D}_0, \mathbf{D}_1)$ between \forall and \exists . If \mathbf{D}_i is one of the models we use \mathbf{D}_{-i} to denote the other model. A position in this game is a pair of sequences $\vec{s}_0 \in D_0^n$ and $\vec{s}_1 \in D_1^n$ with $n \leq k$. The game consists of k rounds where in round $n + 1$ the following steps are made

1. \forall chooses an element d_i in one of the \mathbf{D}_i ,
2. \exists responds with an element d_{-i} in the model \mathbf{D}_{-i} .
3. Let $\vec{s}_i \in D_i^n$ be the sequences of elements chosen up to round n , they are extended to $\vec{s}_i' := \vec{s}_i \cdot d_i$. Player \exists survives the round iff she does not get stuck and the function $f_{n+1} : \vec{s}_0' \mapsto \vec{s}_1'$ is a partial isomorphism of one-step models.

Player \exists wins iff she can survive all k rounds. Given $n \leq k$ and $\vec{s}_i \in D_i^n$ such that $f_n : \vec{s}_0 \mapsto \vec{s}_1$ is a partial isomorphism, we use $\text{EF}_k^\perp(\mathbf{D}_0, \mathbf{D}_1)@(\vec{s}_0, \vec{s}_1)$ to denote the (initialized) game where n moves have been played and $k - n$ moves are left to be played.

Lemma 3.11 *The following are equivalent*

1. $\mathbf{D}_0 \equiv_k^{\text{FOE}} \mathbf{D}_1$,
2. $\mathbf{D}_0 \sim_k^\perp \mathbf{D}_1$,
3. \exists has a winning strategy in $\text{EF}_k^\perp(\mathbf{D}_0, \mathbf{D}_1)$.

Proof. Step (1) to (2) is direct by Proposition 3.9. For (2) to (3) we give a winning strategy for \exists in $\text{EF}_k^\perp(\mathbf{D}_0, \mathbf{D}_1)$. We do it by showing the following claim

Claim 1 Let $\mathbf{D}_0 \sim_k^\perp \mathbf{D}_1$ and $\vec{s}_i \in D_i^n$ be such that $n < k$ and $f_n : \vec{s}_0 \mapsto \vec{s}_1$ is a partial isomorphism; then \exists can survive one more round in $\text{EF}_k^\perp(\mathbf{D}_0, \mathbf{D}_1)@(\vec{s}_0, \vec{s}_1)$.

PROOF OF CLAIM. Let \forall pick $d_i \in D_i$ such that the type of d_i is $T \subseteq A$. If d_i had already been played then \exists picks the same element as before and $f_{n+1} = f_n$. If d_i is new and $|T|_{\mathbf{D}_i} \geq k$ then, as at most $n < k$ elements have been played, there is always some new $d_{-i} \in D_{-i}$ that \exists can choose that matches d_i . If $|T|_{\mathbf{D}_i} = m < k$ then we know that $|T|_{\mathbf{D}_{-i}} = m$. Therefore, as d_i is new and f_n is injective, there must be a $d_{-i} \in D_{-i}$ that \exists can choose. \blacktriangleleft

Step (3) to (1) is a standard result [3, Corollary 2.2.9] which we prove anyway because we will need to extend it later. We prove the following loaded statement

Claim 2 Let $\vec{s}_i \in D_i^n$ and $\varphi(z_1, \dots, z_n) \in \text{FOE}_1(A)$ be such that $qr(\varphi) \leq k - n$. If \exists has a winning strategy in $\text{EF}_k^\perp(\mathbf{D}_0, \mathbf{D}_1)@(\vec{s}_0, \vec{s}_1)$ then $\mathbf{D}_0 \models \varphi(\vec{s}_0)$ iff $\mathbf{D}_1 \models \varphi(\vec{s}_1)$.

PROOF OF CLAIM. If φ is atomic the claim holds because of $f_n : \vec{s}_0 \mapsto \vec{s}_1$ being a partial isomorphism. Boolean cases are straightforward. Let $\varphi(z_1, \dots, z_n) = \exists x. \psi(z_1, \dots, z_n, x)$ and suppose $\mathbf{D}_0 \models \varphi(\vec{s}_0)$. Hence, there exists $d_0 \in D_0$ such that $\mathbf{D}_0 \models \psi(\vec{s}_0, d_0)$. By hypothesis we know that \exists has a winning strategy for $\text{EF}_k^\perp(\mathbf{D}_0, \mathbf{D}_1)@(\vec{s}_0, \vec{s}_1)$. Therefore,

if \forall picks $d_0 \in D_0$ she can respond with some $d_1 \in D_1$ and has a winning strategy for $\text{EF}_k^=(\mathbf{D}_0, \mathbf{D}_1) @ (\vec{s}_0 \cdot d_0, \vec{s}_1 \cdot d_1)$. By induction hypothesis, because $qr(\psi) \leq k - (n + 1)$, we have that $\mathbf{D}_0 \models \psi(\vec{s}_0, d_0)$ iff $\mathbf{D}_1 \models \psi(\vec{s}_1, d_1)$ and hence $\mathbf{D}_1 \models \exists x. \psi(\vec{s}_1, x)$. The other direction is symmetric. \blacktriangleleft

Combining these claims finishes the proof of the lemma. QED

Theorem 3.12 *Every $\psi \in \text{FOE}_1(A)$ is equivalent to a formula in basic form.*

Proof. Let $qr(\psi) = k$ and let $\llbracket \psi \rrbracket$ be the models satisfying ψ . As $\mathfrak{M}_1 / \equiv_k^{\text{FOE}}$ is the same as $\mathfrak{M}_1 / \sim_k^-$ by Lemma 3.11, it is easy to see that $\psi \equiv \bigvee \{ \varphi_E^- \mid E \in \llbracket \psi \rrbracket / \sim_k^- \}$. Now it only remains to see that each φ_E^- is equivalent to $\nabla_{\text{FOE}}(\vec{T}, \Pi)$ for some $\Pi \subseteq \wp A$ and $T_i \subseteq A$.

The crucial observation is that we will use \vec{T} and Π to give a specification of the types “element by element”. Let $\mathbf{D} \in E$ be a representative. Call $S_1, \dots, S_n \subseteq A$ to the types such that $|S_i|_{\mathbf{D}} = n_i < k$ and $S'_1, \dots, S'_m \subseteq A$ to those satisfying $|S'_i|_{\mathbf{D}} \geq k$. The size of the sequence \vec{T} is defined to be $(\sum_{i=1}^n n_i) + k \times m$ where \vec{T} contains exactly n_i occurrences of type T_i and k occurrences of each S'_j . On the other hand $\Pi = \{S'_1, \dots, S'_m\}$. It is straightforward to check that φ_E^- is equivalent to $\nabla_{\text{FOE}}(\vec{T}, \Pi)$, however, the quantifier rank of the latter is only bounded by $k \times 2^{|A|} + 1$. QED

3.1.3 Normal form for FOE_1^∞

The logic FOE_1^∞ basically extends FOE_1 with the capacity to tear apart finite and infinite sets of elements. This is reflected in the normal form for FOE_1^∞ by adding extra constraints to the normal form of FOE_1 .

Definition 3.13 A formula $\varphi \in \text{FOE}_1^\infty(A)$ is in *basic form* if $\varphi = \bigvee \nabla_{\text{FOE}^\infty}(\vec{T}, \Pi, \Sigma)$ where each disjunct is of the form

$$\nabla_{\text{FOE}^\infty}(\vec{T}, \Pi, \Sigma) = \nabla_{\text{FOE}}(\vec{T}, \Pi \cup \Sigma) \wedge \nabla_\infty(\Sigma)$$

where

$$\nabla_\infty(\Sigma) := \bigwedge_{S \in \Sigma} \exists^\infty y. \tau_S(y) \wedge \forall^\infty y. \bigvee_{S \in \Sigma} \tau_S(y)$$

for some set of types $\Pi, \Sigma \subseteq \wp A$ and each $T_i \subseteq A$.

A small argument reveals that, intuitively, every disjunct expresses that each one-step model satisfying it admits a partition of its domain in three parts:

- (i) Distinct elements t_1, \dots, t_n with type T_1, \dots, T_n ,
- (ii) Finitely many elements whose types belong to Π , and
- (iii) For each $S \in \Sigma$, infinitely many elements with type S .

In the same way as before, we define a relation \sim_k^∞ which can be clearly seen to be a refinement of \sim_k^- which adds information about the (in-)finiteness of the types.

Definition 3.14 Let \mathbf{D} and \mathbf{D}' be one-step models. For every $k \in \mathbb{N}$ we define

$$\begin{aligned} \mathbf{D} \sim_0^\infty \mathbf{D}' &\iff \text{always} \\ \mathbf{D} \sim_{k+1}^\infty \mathbf{D}' &\iff \mathbf{D} \sim_{k+1}^\infty \mathbf{D}' \text{ and} \\ &\quad \forall S \subseteq A \left(|S|_{\mathbf{D}}, |S|_{\mathbf{D}'} < \omega \text{ or } |S|_{\mathbf{D}}, |S|_{\mathbf{D}'} \geq \omega \right) \end{aligned}$$

Proposition 3.15 *The following hold*

- (i) \sim_k^∞ is an equivalence relation,
- (ii) \sim_k^∞ has finite index,
- (iii) \sim_k^∞ is a refinement of \sim_k^∞ ,
- (iv) Every $E \in \mathfrak{M}_1 / \sim_k^\infty$ is characterized by a formula $\varphi_E^\infty \in \text{FOE}_1^\infty(A)$ with $qr(\varphi) = k$.

Proof. We only prove the last point, for $k > 0$. Let $E \in \mathfrak{M}_1 / \sim_k^\infty$ and let $\mathbf{D} \in E$ be a representative of the class. Let $E' \in \mathfrak{M}_1 / \sim_k^\infty$ be the equivalence class of \mathbf{D} with respect to \sim_k^∞ . Call $S_1, \dots, S_n \subseteq A$ to the types such that $|S_i|_{\mathbf{D}} \geq \omega$ and $S'_1, \dots, S'_m \subseteq A$ to those satisfying $|S'_i|_{\mathbf{D}} < \omega$. Now define

$$\varphi_E^\infty := \varphi_{E'}^\infty \wedge \bigwedge_{i \leq n} \exists^\infty x. \tau_{S_i}(x) \wedge \bigwedge_{i \leq m} \forall^\infty x. \neg \tau_{S'_i}(x).$$

It is easy to see that $qr(\varphi_E^\infty) = k$ and that $\mathbf{D}' \models \varphi_E^\infty$ iff $\mathbf{D}' \in E$.

QED

Now we give a notion of Ehrenfeucht-Fraïssé game for FOE_1^∞ . In this case the game extends EF_k^∞ with a move for \exists^∞ .

Definition 3.16 Let $\mathbf{D}_0 = (D_0, V_0)$ and $\mathbf{D}_1 = (D_1, V_1)$ be one-step models. We define the game $\text{EF}_k^\infty(\mathbf{D}_0, \mathbf{D}_1)$ between \forall and \exists . A position in this game is a pair of sequences $\vec{s}_0 \in D_0^n$ and $\vec{s}_1 \in D_1^n$ with $n \leq k$. The game consists of k rounds where in round $n+1$ the following steps are made. First \forall chooses to perform one of the following types of moves:

- (a) Second-order move
 1. \forall chooses an infinite set $X_i \subseteq D_i$,
 2. \exists responds with an infinite set $X_{-i} \subseteq D_{-i}$,
 3. \forall chooses an element $x_{-i} \in X_{-i}$,
 4. \exists responds with an element $x_i \in X_i$.

- (b) First-order move
 1. \forall chooses an element $d_i \in D_i$,
 2. \exists responds with an element $d_{-i} \in D_{-i}$.

Let $\vec{s}_i \in D_i^n$ be the sequences of elements chosen up to round n , they are extended to $\vec{s}_i' := \vec{s}_i \cdot d_i$. \exists survives the round iff she does not get stuck and the function $f_{n+1} : \vec{s}_0' \mapsto \vec{s}_1'$ is a partial isomorphism of one-step models.

This game can be seen as an adaptation of the Ehrenfeucht-Fraïssé game for monotone generalized quantifiers found in [11] to the case of full monadic first-order logic.

Lemma 3.17 *The following are equivalent*

1. $\mathbf{D}_0 \equiv_k^{\text{FOE}^\infty} \mathbf{D}_1$,
2. $\mathbf{D}_0 \sim_k^\infty \mathbf{D}_1$,
3. \exists has a winning strategy in $\text{EF}_k^\infty(\mathbf{D}_0, \mathbf{D}_1)$.

Proof. Step (1) to (2) is direct by Proposition 3.15. For (2) to (3) we show

Claim 1 Let $\mathbf{D}_0 \sim_k^\infty \mathbf{D}_1$ and $\vec{s}_i \in D_i^n$ be such that $n < k$ and $f_n : \vec{s}_0 \mapsto \vec{s}_1$ is a partial isomorphism; then \exists can survive one more round in $\text{EF}_k^\infty(\mathbf{D}_0, \mathbf{D}_1) @ (\vec{s}_0, \vec{s}_1)$.

PROOF OF CLAIM. We focus on the second-order moves because the first-order moves are the same as in the corresponding Claim of Lemma 3.11. Let \forall choose an infinite set $X_i \subseteq D_i$, we would like \exists to choose a set $X_{-i} \subseteq D_{-i}$ such that the following conditions hold:

- (a) Let $\vec{r}_j := \vec{s}_j \cap X_j$ be the restriction of \vec{s}_j to the elements of X_j . We want $f' : \vec{r}_0 \mapsto \vec{r}_1$ to be a well-defined partial isomorphism between the restricted models $\mathbf{D}_0|X_0$ and $\mathbf{D}_1|X_1$,
- (b) $\mathbf{D}_0|X_0 \sim_{m+1}^\infty \mathbf{D}_1|X_1$ where $m = |r_i|$. That is, for all $S \subseteq A$,
 - (i) if $|S|_{X_i} < m + 1$ then $|S|_{X_{-i}} = |S|_{X_i}$,
 - (ii) if $|S|_{X_i} \geq m + 1$ then $|S|_{X_{-i}} \geq m + 1$,
- (c) X_{-i} is infinite.

First we prove that such a set exists and then we will use it to prove the claim.

To satisfy item (a) she just needs to add to X_{-i} the elements connected to X_i by f_n ; this is not a problem. For item (b) the only hinder is that, while adding elements to X_{-i} , we cannot make use of those elements of D_{-i} which are in \vec{s}_{-i} but not in \vec{r}_{-i} , since that would break condition (a). It is not difficult to see that because $m + 1 \leq k$ and $\mathbf{D}_0 \sim_k^\infty \mathbf{D}_1$ implies $\mathbf{D}_0 \sim_k^\infty \mathbf{D}_1$ then we always have enough elements to satisfy the required constraints. For the last item observe that as X_i is infinite but there are only finitely many types, hence there must be some S such that $|S|_{X_i} \geq \omega$. It is then safe to add infinitely many elements for S in X_{-i} while considering point (b). Moreover, the existence of infinitely many elements satisfying S in D_{-i} are guaranteed by $\mathbf{D}_0 \sim_k^\infty \mathbf{D}_1$.

Having shown that \exists can choose a set X_{-i} satisfying the above conditions, we continue the proof of the claim as follows: as $\mathbf{D}_0|X_0 \sim_{m+1}^\infty \mathbf{D}_1|X_1$ and $f' : \vec{r}_0 \mapsto \vec{r}_1$ is a partial isomorphism between them (with sequences of length m) then by the first claim of Lemma 3.11, we know that \exists can survive one round in $\text{EF}_{m+1}^\infty(\mathbf{D}_0|X_0, \mathbf{D}_1|X_1) @ (\vec{r}_0, \vec{r}_1)$. In particular, she can survive the “first-order part” of the second-order move we were considering. This finishes the proof of the claim. \blacktriangleleft

Going back to the proof of the lemma, for step (3) to (1) we prove the following claim.

Claim 2 Let $\vec{s}_i \in D_i^n$ and $\varphi(z_1, \dots, z_n) \in \text{FOE}_1^\infty(A)$ be such that $qr(\varphi) \leq k - n$. If \exists has a winning strategy in $\text{EF}_k^\infty(\mathbf{D}_0, \mathbf{D}_1) @ (\vec{s}_0, \vec{s}_1)$ then $\mathbf{D}_0 \models \varphi(\vec{s}_0)$ iff $\mathbf{D}_1 \models \varphi(\vec{s}_1)$.

PROOF OF CLAIM. All the cases involving operators of FOE_1 are the same as in Lemma 3.11. We prove the inductive case for the generalized quantifier. Let $\varphi(z_1, \dots, z_n)$ be of the form $\exists^\infty x. \psi(z_1, \dots, z_n, x)$ and let $\mathbf{D}_0 \models \varphi(\vec{s}_0)$. Hence, there is an *infinite* set $X_0 \subseteq D_0$ such that

$$\mathbf{D}_0 \models \psi(\vec{s}_0, x_0) \text{ if and only if } x_0 \in X_0. \quad (7)$$

By hypothesis we know that \exists has a winning strategy for $\text{EF}_k^\infty(\mathbf{D}_0, \mathbf{D}_1) @ (\vec{s}_0, \vec{s}_1)$. Therefore, if \forall plays a second-order move by picking $X_0 \subseteq D_0$ she can respond with some infinite set $X_1 \subseteq D_1$. We claim that $\mathbf{D}_1 \models \psi(\vec{s}_1, x_1)$ for every $x_1 \in X_1$. First observe that if this holds then $X'_1 := \{d_1 \in D_1 \mid \mathbf{D}_1 \models \psi(\vec{s}_1, d_1)\}$ must be infinite, and hence $\mathbf{D}_1 \models \exists^\infty x. \psi(\vec{s}_1, x)$.

Assume, for a contradiction, that $\mathbf{D}_1 \not\models \psi(\vec{s}_1, x'_1)$ for some $x'_1 \in X_1$. Let \forall play that x'_1 as the second part of his move. Then, as \exists has a winning strategy, she will respond with some $x'_0 \in X_0$ such that she has a winning strategy for $\text{EF}_k^\infty(\mathbf{D}_0, \mathbf{D}_1) @ (\vec{s}_0 \cdot x'_0, \vec{s}_1 \cdot x'_1)$. By induction hypothesis, as $qr(\psi) \leq k - (n + 1)$, this means that $\mathbf{D}_0 \models \psi(\vec{s}_0, x'_0)$ iff $\mathbf{D}_1 \models \varphi(\vec{s}_1, x'_1)$ which contradicts (7). The other direction is symmetric. \blacktriangleleft

Combining the claims finishes the proof of the lemma. QED

Theorem 3.18 Every $\varphi \in \text{FOE}_1^\infty(A)$ is equivalent to a formula in basic form.

Proof. This can be proved using the same technique as in Theorem 3.12. Hence we only focus on showing that $\varphi_E^\infty \equiv \nabla_{\text{FOE}^\infty}(\vec{T}, \Pi, \Sigma)$ for some $\Pi, \Sigma \subseteq \wp A$ and $T_i \subseteq A$. Recall that

$$\varphi_E^\infty = \varphi_{E'}^\infty \wedge \bigwedge_{i \leq n} \exists^\infty x. \tau_{S_i}(x) \wedge \bigwedge_{i \leq m} \forall^\infty x. \neg \tau_{S'_i}(x)$$

and from Theorem 3.12 we know that this is equivalent to

$$\varphi_E^\infty = \nabla_{\text{FOE}}(\vec{T}', \Pi') \wedge \bigwedge_{i \leq n} \exists^\infty x. \tau_{S_i}(x) \wedge \bigwedge_{i \leq m} \forall^\infty x. \neg \tau_{S'_i}(x)$$

for some $\Pi' \subseteq \wp A$ and $T'_i \subseteq A$. Now separate Π' as $\Pi' = \Pi \uplus \Sigma$ where Π is composed of the finite types (i.e., the S'_i) and Σ is composed of the infinite types (i.e., the S_i). We claim that $\varphi_E^\infty \equiv \nabla_{\text{FOE}^\infty}(\vec{T}', \Pi, \Sigma)$. After a minor rewritting, we have to show that the following two formulas are equivalent:

$$\nabla_{\text{FOE}}(\vec{T}', \Pi') \wedge \bigwedge_{S \in \Sigma} \exists^\infty x. \tau_S(x) \wedge \bigwedge_{S \in \Pi} \forall^\infty x. \neg \tau_S(x) \quad (8)$$

$$\nabla_{\text{FOE}}(\vec{T}', \Pi \cup \Sigma) \wedge \bigwedge_{S \in \Sigma} \exists^\infty y. \tau_S(y) \wedge \bigvee_{S \in \Sigma} \tau_S(y) \quad (9)$$

It is clear that $\nabla_{\text{FOE}}(\vec{T}', \Pi')$ is equivalent to $\nabla_{\text{FOE}}(\vec{T}', \Pi \cup \Sigma)$ and that the infinite-existential parts are also equivalent. Therefore we focus on the last part of both formulas.

\Rightarrow We show that (8) implies $\forall^\infty y. \bigvee_{S \in \Sigma} \tau_S(y)$. By contrapositive assume that there are infinitely many elements satisfying $\bigwedge_{S \in \Sigma} \neg \tau_S(y)$, that is, infinitely many elements with their type outside Σ . Because of the FOE part, these elements should have their type in $\Pi' = \Pi \cup \Sigma$. Therefore, as the number of types is finite, there must be some type $P \in \Pi$ such that $\exists^\infty x. \tau_P(x)$ which contradicts the last part of (8).

\Leftarrow We show that (9) implies $\bigwedge_{S \in \Pi} \forall^\infty x. \neg \tau_S(x)$. By contrapositive assume there is some type $P \in \Pi$ such that $\exists^\infty x. \tau_P(x)$. Then, in particular, these elements do not have their type in Σ and hence there are infinitely many elements satisfying $\bigwedge_{S \in \Sigma} \neg \tau_S(y)$ which contradicts the last part of (9). QED

3.2 One-step monotonicity

Given a one-step logic $\mathcal{L}_1(A)$ and formula $\varphi \in \mathcal{L}_1(A)$. We say that φ is *monotone in* $a \in A$ if for every one step model $(D, V : A \rightarrow \wp D)$ and assignment $g : \text{iVar} \rightarrow D$,

$$\text{If } (D, V), g \models \varphi \text{ and } V(a) \subseteq E \text{ then } (D, V[a \mapsto E]), g \models \varphi.$$

We use $\mathcal{L}_1^+(A)$ to denote the fragment of $\mathcal{L}_1(A)$ composed of formulas monotone in all $a \in A$.

Monotonicity is usually tightly related to positivity. If the quantifiers are well-behaved (i.e., monotone) then a formula φ will usually be monotone in $a \in A$ iff a has positive polarity in φ , that is, if it occurs under an even number of negations. This is the case for all one-step logics considered in this article. In this section we give a syntactic characterization of monotonicity for several one-step logics.

Convention. Given $S \cup \{a\} \subseteq A$ we use τ_S^a to denote the *a-positive A-type* defined as

$$\tau_S^a(x) := \bigwedge_{b \in S} b(x) \wedge \bigwedge_{b \in A \setminus (S \cup \{a\})} \neg b(x).$$

Intuitively, τ_S^a is an *A-type* where a can only occur positively and all other predicates occur either positively or negatively.

3.2.1 Monotone fragment of FO_1

Theorem 3.19 *A formula of $\text{FO}_1(A)$ is monotone in $a \in A$ iff it is equivalent to a sentence given by*

$$\varphi ::= \psi \mid a(x) \mid \exists x. \varphi(x) \mid \forall x. \varphi(x) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$$

where $\psi \in \text{FO}_1(A \setminus \{a\})$. We denote this fragment as $\text{FO}_1\text{M}_a(A)$.

The result will follow from the following two lemmas.

Lemma 3.20 *If $\varphi \in \text{FO}_1\text{M}_a(A)$ then φ is monotone in a .*

Proof. The proof is a routine argument by induction on the complexity of φ . QED

To prove that the fragment is complete we need to show that every formula which is monotone in a is equivalent to some formula in FO_1M_a . We prove a stronger result: we give a translation that constructively maps arbitrary formulas into FO_1M_a . The interesting observation is that the translation will preserve truth iff the given formula is monotone in a .

Lemma 3.21 *There exists a translation $(-)^{\circ} : \text{FO}_1(A) \rightarrow \text{FO}_1\text{M}_a(A)$ such that a formula $\varphi \in \text{FO}_1(A)$ is monotone in a if and only if $\varphi \equiv \varphi^{\circ}$.*

Proof. To define the translation we assume that φ is in a normal form $\bigvee \nabla_{\text{FO}}(\Sigma)$ where

$$\nabla_{\text{FO}}(\Sigma) := \bigwedge_{S \in \Sigma} \exists x. \tau_S(x) \wedge \forall x. \bigvee_{S \in \Sigma} \tau_S(x)$$

for some types $\Sigma \subseteq \wp A$. For the translation, let $(\bigvee \nabla_{\text{FO}}(\Sigma))^{\circ} := \bigvee \nabla_{\text{FO}}^a(\Sigma)$ and

$$\nabla_{\text{FO}}^a(\Sigma) := \bigwedge_{S \in \Sigma} \exists x. \tau_S^a(x) \wedge \forall x. \bigvee_{S \in \Sigma} \tau_S^a(x)$$

From the construction it is clear that $\varphi^{\circ} \in \text{FO}_1\text{M}_a(A)$ and therefore the right-to-left direction of the lemma is immediate by Lemma 3.20. For the left-to-right direction assume that φ is monotone in a , we have to prove that $(D, V) \models \varphi$ iff $(D, V) \models \varphi^{\circ}$.

\Rightarrow This direction is trivial.

\Leftarrow Assume $(D, V) \models \varphi^{\circ}$ and let $\nabla_{\text{FO}}^a(\Sigma)$ be such that $(D, V) \models \nabla_{\text{FO}}^a(\Sigma)$. It is safe to assume that the *only* (a -positive) types realized in (D, V) are exactly those in Σ and that every type has a (single) distinct witness. For every $S \in \Sigma$, let d_S be the witness of the a -positive type $\tau_S^a(x)$ in (D, V) . Let $U := \{d_S \mid S \in \Sigma, a \notin S\}$ and $V' := V[a \mapsto V(a) \setminus U]$.

Claim 1 $(D, V') \models \nabla_{\text{FO}}(\Sigma)$.

PROOF OF CLAIM. First we show that the existential part of the normal form is satisfied. That is, that for every $S \in \Sigma$ we have a witness for the *full* type $\tau_S(x)$. If $a \in S$ the witness is given by φ° , that is, d_S . If $a \notin S$ then we specially crafted d_S to be a witness. The universal part is clearly satisfied. \blacktriangleleft

To finish it is enough to observe that, by monotonicity of φ , we get $(D, V) \models \varphi$. QED

Putting together the above lemmas we obtain Theorem 3.19. Moreover, a careful analysis of the translation gives us the following corollary, providing normal forms for the monotone fragment of FO_1 .

Corollary 3.22 *Let $\varphi \in \text{FO}_1(A)$, the following hold:*

- (i) *The formula φ is monotone in $a \in A$ iff it is equivalent to a formula in the basic form $\bigvee \nabla_{\text{FO}}^a(\Sigma)$ for some types $\Sigma \subseteq \wp A$.*
- (ii) *The formula φ is monotone in every $a \in A$ (i.e., $\varphi \in \text{FO}_1^+(A)$) iff φ is equivalent to a formula in the basic form $\bigvee \nabla_{\text{FO}}^+(\Sigma)$ for some types $\Sigma \subseteq \wp A$, where*

$$\nabla_{\text{FO}}^+(\Sigma) := \bigwedge_{S \in \Sigma} \exists x. \tau_S^+(x) \wedge \forall x. \bigvee_{S \in \Sigma} \tau_S^+(x).$$

3.2.2 Monotone fragment of FOE_1^∞

Theorem 3.23 *A formula of $\text{FOE}_1^\infty(A)$ is monotone in $a \in A$ iff it is equivalent to a sentence given by*

$$\varphi ::= \psi \mid a(x) \mid \exists x.\varphi(x) \mid \forall x.\varphi(x) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists^\infty x.\varphi(x) \mid \forall^\infty x.\varphi(x)$$

where $\psi \in \text{FOE}_1^\infty(A \setminus \{a\})$. We denote this fragment as $\text{FOE}_1^\infty M_a(A)$.

Observe that $x \approx y$ and $x \not\approx y$ are included in the case $\psi \in \text{FOE}_1^\infty(A \setminus \{a\})$. The result will follow from the following two lemmas.

Lemma 3.24 *If $\varphi \in \text{FOE}_1^\infty M_a(A)$ then φ is monotone in a .*

Proof. The proof is basically the same as Lemma 3.20. That is, we show by induction, that any one-step formula φ in the fragment (which may not be a sentence) satisfies, for every one-step model $(D, V : A \rightarrow \wp D)$ and assignment $g : \text{iVar} \rightarrow D$,

$$\text{If } (D, V), g \models \varphi \text{ and } V(a) \subseteq E \text{ then } (D, V[a \mapsto E]), g \models \varphi.$$

We focus on the cases for the generalized quantifiers. Let $(D, V), g \models \varphi$ and $V(a) \subseteq E$.

- Case $\varphi = \exists^\infty x.\varphi'(x)$. By definition there exists an infinite set $I \subseteq D$ such that for all $d \in I$ we have $(D, V), g[x \mapsto d] \models \varphi'(x)$. By induction hypothesis $(D, V[a \mapsto E]), g[x \mapsto d] \models \varphi'(x)$ for all $d \in I$. Therefore $(D, V[a \mapsto E]), g \models \exists^\infty x.\varphi'(x)$.
- Case $\varphi = \forall^\infty x.\varphi'(x)$. By definition there exists $I \subseteq D$ such that for all $d \in I$ we have $(D, V), g[x \mapsto d] \models \varphi'(x)$ and $D \setminus I$ is finite. By induction hypothesis $(D, V[a \mapsto E]), g[x \mapsto d] \models \varphi'(x)$ for all $d \in I$. Therefore $(D, V[a \mapsto E]), g \models \forall^\infty x.\varphi'(x)$.

QED

Lemma 3.25 *There exists a translation $(-)^o : \text{FOE}_1^\infty(A) \rightarrow \text{FOE}_1^\infty M_a(A)$ such that a formula $\varphi \in \text{FOE}_1^\infty(A)$ is monotone in a if and only if $\varphi \equiv \varphi^o$.*

Proof. We assume that φ is in a normal form $\bigvee \nabla_{\text{FOE}^\infty}(\vec{T}, \Pi, \Sigma)$ where

$$\nabla_{\text{FOE}^\infty}(\vec{T}, \Pi, \Sigma) = \nabla_{\text{FOE}}(\vec{T}, \Pi \cup \Sigma) \wedge \nabla_\infty(\Sigma).$$

for some sets of types $\Pi, \Sigma \subseteq \wp A$ and each $T_i \subseteq A$. For the translation, we define

$$(\bigvee \nabla_{\text{FOE}^\infty}(\vec{T}, \Pi, \Sigma))^o := \bigvee \nabla_{\text{FOE}^\infty}^a(\vec{T}, \Pi, \Sigma)$$

where

$$\begin{aligned} \nabla_{\text{FOE}^\infty}^a(\vec{T}, \Pi, \Sigma) &:= \nabla_{\text{FOE}}^a(\vec{T}, \Pi \cup \Sigma) \wedge \nabla_\infty^a(\Sigma) \\ \nabla_{\text{FOE}}^a(\vec{T}, \Lambda) &:= \exists \vec{x}. (\text{diff}(\vec{x}) \wedge \bigwedge_i \tau_{T_i}^a(x_i) \wedge \forall z. (\text{diff}(\vec{x}, z) \rightarrow \bigvee_{S \in \Lambda} \tau_S^a(z))) \\ \nabla_\infty^a(\Sigma) &:= \bigwedge_{S \in \Sigma} \exists^\infty y. \tau_S^a(y) \wedge \forall^\infty y. \bigvee_{S \in \Sigma} \tau_S^a(y). \end{aligned}$$

From the construction it is clear that $\varphi^\circ \in \text{FOE}_1^\infty \mathbf{M}_a(A)$ and therefore the right-to-left direction of the lemma is immediate by Lemma 3.24. For the left-to-right direction assume that φ is monotone in a , we have to prove that $(D, V) \models \varphi$ iff $(D, V) \models \varphi^\circ$.

\Rightarrow This direction is trivial.

\Leftarrow Assume $(D, V) \models \varphi^\circ$. Observe that the elements of D can be partitioned as follows:

- (a) Elements $t_i \in D$ witnessing the a -positive type $T_i \in \vec{T}$,
- (b) Disjoint sets $\{D_S \subseteq D \mid S \in \Sigma\}$ such that each D_S is *infinite* and every $d \in D_S$ is a witness for the a -positive type $S \in \Sigma$,
- (c) A *finite* set $D_\Pi \subseteq D$ of witnesses of the a -positive types in Π .

In what follows we assume that every $d \in D$ is the witness of some fixed a -positive type $S_d \in \vec{T} \cup \Pi \cup \Sigma$. We define a new valuation V' as $V'(b) = V(b)$ for all $b \in A \setminus \{a\}$ and $V'(a) := \{d \in D \mid a \in S_d\}$. Observe that $V'(a) \subseteq V(a)$.

Claim 1 $(D, V') \models \varphi$.

PROOF OF CLAIM. First we check that $(D, V') \models \nabla_{\text{FOE}}(\vec{T}, \Pi \cup \Sigma)$. It is easy to see that the elements t_i work as witnesses for the *full* types T_i . That is $(D, V') \models \tau_{T_i}(t_i)$ for every i . To prove the universal part of the formula it is enough to show that

1. Every element $d \in D_\Pi$ realizes the full type $S_d \in \Pi$,
2. For all $S \in \Sigma$, every element of D_S realizes the full type S .

Let d be an element of either D_Π or any of the D_S . By (b) and (c) we know $(D, V) \models \tau_{S_d}^a(d)$. If $a \in S_d$ we can trivially conclude $(D, V') \models \tau_{S_d}(d)$. If $a \notin S_d$, by definition of V' we know that $d \notin V'(a)$ and hence we can also conclude that $(D, V') \models \tau_{S_d}(d)$.

To prove that $(D, V') \models \bigwedge_{S \in \Sigma} \exists^\infty y. \tau_S(y) \wedge \forall^\infty y. \bigvee_{S \in \Sigma} \tau_S(y)$ we only need to observe that the existential part is satisfied because each D_S is infinite by (c) and the universal part is satisfied because the set $D_\Pi \cup \vec{T}$ is finite by (b). \blacktriangleleft

To finish the proof of the lemma, note that by monotonicity of φ we get $(D, V) \models \varphi$. QED

Putting together the above lemmas we obtain Theorem 3.23. Moreover, a careful analysis of the translation gives us the following corollary, providing normal forms for the monotone fragment of FOE_1^∞ .

Corollary 3.26 *Let $\varphi \in \text{FOE}_1^\infty(A)$, the following hold:*

- (i) *The formula φ is monotone in $a \in A$ iff it is equivalent to a formula in the basic form $\bigvee \nabla_{\text{FOE}^\infty}^a(\vec{T}, \Pi, \Sigma)$ for some types $\Pi, \Sigma \subseteq \wp A$ and each $T_i \subseteq A$.*

(ii) The formula φ is monotone in every $a \in A$ (i.e., $\varphi \in \text{FOE}_1^{\infty+}(A)$) iff it is equivalent to a formula in the basic form $\bigvee \nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$ for types $\Pi, \Sigma \subseteq \wp A$ and $T_i \subseteq A$, where

$$\begin{aligned}\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma) &:= \nabla_{\text{FOE}}^+(\vec{T}, \Pi \cup \Sigma) \wedge \nabla_\infty^+(\Sigma) \\ \nabla_{\text{FOE}}^+(\vec{T}, \Lambda) &:= \exists \vec{x}. (\text{diff}(\vec{x}) \wedge \bigwedge_i \tau_{T_i}^+(x_i) \wedge \forall z. (\text{diff}(\vec{x}, z) \rightarrow \bigvee_{S \in \Lambda} \tau_S^+(z))) \\ \nabla_\infty^+(\Sigma) &:= \bigwedge_{S \in \Sigma} \exists^\infty y. \tau_S^+(y) \wedge \forall^\infty y. \bigvee_{S \in \Sigma} \tau_S^+(y).\end{aligned}$$

3.3 One-step continuity

Consider a one-step logic $\mathcal{L}_1(A)$ and formula $\varphi \in \mathcal{L}_1(A)$. We say that φ is *continuous* in $a \in A$ if φ is monotone in a and additionally, for every (D, V) and assignment $g : \text{iVar} \rightarrow D$,

$$\text{if } (D, V), g \models \varphi \text{ then } \exists U \subseteq_\omega V(a) \text{ such that } (D, V[a \mapsto U]), g \models \varphi.$$

It will be useful to give a syntactic characterization of continuity for several one-step logics.

3.3.1 Continuous fragment of FO_1

Theorem 3.27 *A formula of $\text{FO}_1(A)$ is continuous in $a \in A$ iff it is equivalent to a sentence given by*

$$\varphi ::= \psi \mid a(x) \mid \exists x. \varphi(x) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$$

where $\psi \in \text{FO}_1(A \setminus \{a\})$. We denote this fragment as $\text{FO}_1\text{C}_a(A)$.

The theorem will follow from the next two lemmas.

Lemma 3.28 *If $\varphi \in \text{FO}_1\text{C}_a(A)$ then φ is continuous in a .*

Proof. First observe that φ is monotone in a by Theorem 3.19. We show, by induction, that any one-step formula φ in the fragment (which may not be a sentence) satisfies, for every one-step model $(D, V : A \rightarrow \wp D)$, assignment $g : \text{iVar} \rightarrow D$,

$$\text{If } (D, V), g \models \varphi \text{ then } \exists U \subseteq_\omega V(a) \text{ such that } (D, V[a \mapsto U]), g \models \varphi.$$

- If $\varphi = \psi \in \text{FO}_1(A \setminus \{a\})$ changes in the a part of the valuation will make no difference and hence the condition is trivial.
- Case $\varphi = a(x)$: if $(D, V), g \models a(x)$ then $g(x) \in V(a)$. Clearly, $g(x) \in V[a \mapsto \{g(x)\}](a)$ and hence $(D, V[a \mapsto \{g(x)\}]), g \models a(x)$.
- Case $\varphi = \varphi_1 \vee \varphi_2$: assume $(D, V), g \models \varphi$. Without loss of generality we can assume that $(D, V), g \models \varphi_1$ and hence by induction hypothesis we have that there is $U \subseteq_\omega V(a)$ such that $(D, V[a \mapsto U]), g \models \varphi_1$ which clearly satisfies $(D, V[a \mapsto U]), g \models \varphi$.
- Case $\varphi = \varphi_1 \wedge \varphi_2$: assume $(D, V), g \models \varphi$. By induction hypothesis we have $U_1, U_2 \subseteq_\omega V(a)$ such that $(D, V[a \mapsto U_1]), g \models \varphi_1$ and $(D, V[a \mapsto U_2]), g \models \varphi_2$. By monotonicity this also holds with $V[a \mapsto U_1 \cup U_2]$ and therefore $(D, V[a \mapsto U_1 \cup U_2]), g \models \varphi$.

- Case $\varphi = \exists x.\varphi'(x)$ and $(D, V), g \models \varphi$. By definition there exists $d \in D$ such that $(D, V), g[x \mapsto d] \models \varphi'(x)$. By induction hypothesis there exists $U \subseteq_\omega V(a)$ such that $(D, V[a \mapsto U]), g[x \mapsto d] \models \varphi'(x)$ and hence $(D, V[a \mapsto U]), g \models \exists x.\varphi'(x)$.

QED

Lemma 3.29 *There exists a translation $(-)^{\Delta} : \text{FO}_1\text{M}_a(A) \rightarrow \text{FO}_1\text{C}_a(A)$ such that a formula $\varphi \in \text{FO}_1\text{M}_a(A)$ is continuous in a if and only if $\varphi \equiv \varphi^{\Delta}$.*

Proof. To define the translation we assume that φ is in basic form $\bigvee \nabla_{\text{FO}}^a(\Sigma)$ where

$$\nabla_{\text{FO}}^a(\Sigma) := \bigwedge_{S \in \Sigma} \exists x.\tau_S^a(x) \wedge \forall x. \bigvee_{S \in \Sigma} \tau_S^a(x)$$

for some types $\Sigma \subseteq \wp A$. For the translation, let $(\bigvee \nabla_{\text{FO}}^a(\Sigma))^{\Delta} := \bigvee \nabla_{\text{FO}}^a(\Sigma)^{\Delta}$ and

$$\nabla_{\text{FO}}^a(\Sigma)^{\Delta} := \bigwedge_{S \in \Sigma} \exists x.\tau_S^a(x) \wedge \forall x. \bigvee_{S \in \Sigma_a^-} \tau_S^a(x)$$

where $\Sigma_a^- := \{S \in \Sigma \mid a \notin S\}$.

From the construction it is clear that $\varphi^{\Delta} \in \text{FO}_1\text{C}_a(A)$ and therefore the right-to-left direction of the lemma is immediate by Lemma 3.28. For the left-to-right direction assume that φ is continuous in a , we have to prove that $(D, V) \models \varphi$ iff $(D, V) \models \varphi^{\Delta}$, for every one-step model (D, V) . We will take a slightly different but equivalent approach.

It is easy to prove that $(D, V) \equiv_{\text{FO}} (D \times \omega, V_{\pi})$ where $D \times \omega$ has countably many copies of each element in D and $V_{\pi}(a) := \{(d, k) \mid d \in V(a), k \in \omega\}$. Moreover, as φ is continuous in a there is $U \subseteq_\omega V_{\pi}(a)$ such that $V'_{\pi} := V[a \mapsto U]$ satisfies $(D \times \omega, V_{\pi}) \models \varphi$ iff $(D \times \omega, V'_{\pi}) \models \varphi$. Therefore, it is enough to prove that $(D \times \omega, V'_{\pi}) \models \varphi$ iff $(D \times \omega, V'_{\pi}) \models \varphi^{\Delta}$.

\Rightarrow Let $(D \times \omega, V'_{\pi}) \models \nabla_{\text{FO}}(\Sigma)$, we show that $(D \times \omega, V'_{\pi}) \models \nabla_{\text{FO}}(\Sigma)^{\Delta}$. The existential part of $\nabla_{\text{FO}}(\Sigma)^{\Delta}$ is trivially true. We have to show that every element of $(D \times \omega, V'_{\pi})$ realizes a a -positive type in Σ_a^- . Let $(d, k) \in D \times \omega$ and T be the (full) type of (d, k) . If $a \notin T$ then trivially $T \in \Sigma_a^-$ and we are done. Suppose $a \in T$. Observe that in $D \times \omega$ we have infinitely many copies of $d \in D$. However, as $V'_{\pi}(a)$ is finite, there must be some (d, k') with type $T' := T \setminus \{a\}$. For $\nabla_{\text{FO}}(\Sigma)$ to be true we must have $T' \in \Sigma$ and hence $T' \in \Sigma_a^-$. It is easy to see that (d, k) realizes the a -positive type T' .

\Leftarrow Let $(D \times \omega, V'_{\pi}) \models \nabla_{\text{FO}}(\Sigma)^{\Delta}$, we show that $(D \times \omega, V'_{\pi}) \models \nabla_{\text{FO}}(\Sigma)$. The existential part is again trivial. For the universal part just observe that $\Sigma_a^- \subseteq \Sigma$. QED

Putting together the above lemmas we obtain Theorem 3.27. Moreover, a careful analysis of the translation gives us the following corollary, providing normal forms for the continuous fragment of FO_1 .

Corollary 3.30 *Let $\varphi \in \text{FO}_1(A)$, the following hold:*

- (i) The formula φ is continuous in $a \in A$ iff it is equivalent to a formula in the basic form $\bigvee \chi_\Sigma^a$ for some types $\Sigma \subseteq \wp A$, where χ_Σ^a is given by

$$\chi_\Sigma^a := \bigwedge_{S \in \Sigma} \exists x. \tau_S^a(x) \wedge \forall x. \bigvee_{S \in \Sigma_a^-} \tau_S^a(x).$$

- (ii) If φ is monotone in A (i.e., $\varphi \in \text{FO}_1^+(A)$) then φ is continuous in $a \in A$ iff it is equivalent to a formula in the basic form $\bigvee \chi_\Sigma^+$ for some types $\Sigma \subseteq \wp A$, where

$$\chi_\Sigma^+ := \bigwedge_{S \in \Sigma} \exists x. \tau_S^+(x) \wedge \forall x. \bigvee_{S \in \Sigma_a^-} \tau_S^+(x).$$

3.3.2 Continuous fragment of FOE_1^∞

Theorem 3.31 A formula of $\text{FOE}_1^\infty(A)$ is continuous in $a \in A$ iff it is equivalent to a sentence given by

$$\varphi ::= \psi \mid a(x) \mid \exists x. \varphi(x) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{W}x.(\varphi, \psi)$$

where $\psi \in \text{FOE}_1^\infty(A \setminus \{a\})$ and $\mathbf{W}x.(\varphi, \psi) := \forall x.(\varphi(x) \vee \psi(x)) \wedge \forall^\infty x. \psi(x)$. We denote this fragment as $\text{FOE}_1^\infty C_a(A)$.

Universal quantification is usually problematic for preserving continuity because of its potentially infinite nature. However, in the case of $\mathbf{W}x.(\varphi, \psi)$, the combination of both quantifiers ensures that all the elements are covered by $\varphi \vee \psi$ but only *finitely many* are required to make φ true (which contains $a \in A$). This gives no trouble for continuity in a .

The theorem will follow from the next two lemmas.

Lemma 3.32 If $\varphi \in \text{FOE}_1^\infty C_a(A)$ then φ is continuous in a .

Proof. Observe that monotonicity of φ is guaranteed by Theorem 3.23. We show, by induction, that any formula of the fragment (which may not be a sentence) satisfies, for every one-step model $(D, V : A \rightarrow \wp D)$ and assignment $g : \text{iVar} \rightarrow D$,

$$\text{If } (D, V), g \models \varphi \text{ then } \exists U \subseteq_\omega V(a) \text{ such that } (D, V[a \mapsto U]), g \models \varphi.$$

We focus on the inductive case of the new quantifier. Let $\varphi' = \mathbf{W}x.(\varphi, \psi)$, i.e.,

$$\varphi' = \forall x. \underbrace{(\varphi(x) \vee \psi(x))}_{\alpha(x)} \wedge \underbrace{\forall^\infty x. \psi(x)}_{\beta}.$$

Let $(D, V), g \models \varphi'$. By induction hypothesis, for every $g_d := g[x \mapsto d]$ which satisfies $(D, V), g_d \models \alpha(x)$ there is $U_d \subseteq_\omega V(a)$ such that $(D, V[a \mapsto U_d]), g_d \models \alpha(x)$. The crucial observation is that because of β , only finitely many elements of d make $\psi(d)$ false. Let $U := \bigcup \{U_d \mid (D, V), g_d \not\models \psi(x)\}$. Note that U is a finite union of finite sets, hence finite.

Claim 1 Let $V_U := V[a \mapsto U]$, we have $(D, V_U), g \models \varphi'$.

PROOF OF CLAIM. It is clear that $(D, V_U), g \models \beta$ because ψ is a -free. To show that the first conjunct is true we have to show that $(D, V_U), g_d \models \varphi(x) \vee \psi(x)$ for every $d \in D$. We consider two cases: (i) if $(D, V), g_d \models \psi(x)$ we are done, again because ψ is a -free; (ii) if the former is not the case then $U_d \subseteq U$; moreover, we knew that $(D, V[a \mapsto U_d]), g_d \models \alpha(x)$ and by monotonicity of $\alpha(x)$ we can conclude that $(D, V_U), g_d \models \alpha(x)$. \blacktriangleleft

This finishes the proof of the lemma. QED

Lemma 3.33 *There exists a translation $(-)^{\Delta} : \text{FOE}_1^{\infty} M_a(A) \rightarrow \text{FOE}_1^{\infty} C_a(A)$ such that a formula $\varphi \in \text{FOE}_1^{\infty} M_a(A)$ is continuous in a if and only if $\varphi \equiv \varphi^{\Delta}$.*

Proof. We assume that φ is in basic normal form, i.e., $\varphi = \bigvee \nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma)$. For the translation let $(\bigvee \nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma))^{\Delta} := \bigvee \nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma)^{\Delta}$ and

$$\nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma)^{\Delta} := \begin{cases} \perp & \text{if } a \in \bigcup \Sigma \\ \nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma) & \text{otherwise.} \end{cases}$$

First we prove the right-to-left direction of the lemma. By Lemma 3.32 it is enough to show that $\varphi^{\Delta} \in \text{FOE}_1^{\infty} C_a(A)$. We focus on the disjuncts of φ^{Δ} . The interesting case is when $a \notin \bigcup \Sigma$. If we rearrange $\nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma)$ and define the formulas φ', ψ as follows:

$$\begin{aligned} \nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma) \equiv \exists \vec{x}. & \left(\text{diff}(\vec{x}) \wedge \bigwedge_i \tau_{T_i}^a(x_i) \wedge \right. \\ & \underbrace{\forall z. (\neg \text{diff}(\vec{x}, z) \vee \bigvee_{S \in \Pi} \tau_S^a(z))}_{\varphi'(\vec{x}, z)} \vee \underbrace{\bigvee_{S \in \Sigma} \tau_S^a(z)}_{\psi(z)} \Big) \wedge \\ & \underbrace{\forall^{\infty} y. \bigvee_{S \in \Sigma} \tau_S^a(y)}_{\psi(y)} \Big) \wedge \bigwedge_{S \in \Sigma} \exists^{\infty} y. \tau_S^a(y), \end{aligned}$$

then we get that

$$\nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma) \equiv \exists \vec{x}. \left(\text{diff}(\vec{x}) \wedge \bigwedge_i \tau_{T_i}^a(x_i) \wedge \mathbf{W}z. (\varphi'(\vec{x}, z), \psi(z)) \right) \wedge \bigwedge_{S \in \Sigma} \exists^{\infty} y. \tau_S^a(y)$$

which, because $a \notin \bigcup \Sigma$, is in the required fragment.

For the left-to-right direction of the lemma we have to prove that $\varphi \equiv \varphi^{\Delta}$.

$\boxed{\Leftarrow}$ Let $(D, V) \models \varphi^{\Delta}$. The only difference between φ and φ^{Δ} is that some disjuncts may have been replaced by \perp . Therefore this direction is trivial.

$\boxed{\Rightarrow}$ Let $(D, V) \models \varphi$. Because φ is continuous in a we may assume that $V(a)$ is finite. Let $\nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma)$ be a disjunct of φ such that $(D, V) \models \nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma)$. If $a \notin \bigcup \Sigma$ we trivially conclude that $(D, V) \models \varphi^{\Delta}$ because the disjunct remains unchanged. Suppose now that $a \in \bigcup \Sigma$, then there must be some $S \in \Sigma$ with $a \in S$. Because $(D, V) \models \nabla_{\text{FOE}^{\infty}}^a(\vec{T}, \Pi, \Sigma)$ we have, in particular, that $(D, V) \models \exists^{\infty} y. \tau_S^a(y)$ and hence $V(a)$ must be infinite which is absurd. QED

Putting together the above lemmas we obtain Theorem 3.31. Moreover, a careful analysis of the translation gives us the following corollary, providing normal forms for the continuous fragment of FOE_1^∞ .

Corollary 3.34 *Let $\varphi \in \text{FOE}_1^\infty(A)$, the following hold:*

- (i) *The formula φ is continuous in $a \in A$ iff it is equivalent to a formula in the basic form $\bigvee \nabla_{\text{FOE}^\infty}^a(\vec{T}, \Pi, \Sigma)$ for some types $\Pi, \Sigma \subseteq \wp A$ and $T_i \subseteq A$ such that $a \notin \bigcup \Sigma$.*
- (ii) *If φ is monotone in every element of A (i.e., $\varphi \in \text{FOE}_1^{\infty+}(A)$) then φ is continuous in $a \in A$ iff it is equivalent to a formula in the basic form $\bigvee \nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$ for some types $\Pi, \Sigma \subseteq \wp A$ and $T_i \subseteq A$ such that $a \notin \bigcup \Sigma$.*

3.4 One-step co-continuity and Boolean duals

Consider a one-step logic $\mathcal{L}_1(A)$ and formula $\varphi \in \mathcal{L}_1(A)$. We say that φ is *co-continuous* in $a \in A$ if φ is monotone in a and, for all (D, V) and $g : \text{iVar} \rightarrow D$,

$$\text{if } (D, V), g \not\models \varphi \text{ then } \exists U \subseteq_\omega V(a) \text{ such that } (D, V[a \mapsto D \setminus U]), g \not\models \varphi.$$

Observe that, already with the abstract definition of Boolean dual given in Definition 2.5 we can prove the expected relationship between the notions of continuity and co-continuity.

Proposition 3.35 *Let $\varphi \in \mathcal{L}_1(A)$, φ is continuous in $a \in A$ iff φ^δ is co-continuous in a .*

Proof. Let φ be continuous in a , we prove that φ^δ is co-continuous in a :

$$\begin{aligned} (D, V) \not\models \varphi^\delta &\text{ iff } (D, V^c) \models \varphi && \text{(Definition 2.5)} \\ &\text{ iff } \exists U \subseteq_\omega V^c(a) \text{ such that } (D, V^c[a \mapsto U]) \models \varphi && (\varphi \text{ continuous}) \\ &\text{ iff } \exists U' \subseteq_\omega V(a) \text{ such that } (D, V[a \mapsto D \setminus U']) \not\models \varphi && \text{(Definition of } V^c) \end{aligned}$$

The proof of the other direction is analogous. QED

To define a syntactic notion of co-continuity we first give a concrete definition of the dualization operator of Definition 2.5 and then show that the one-step language FOE_1^∞ is closed under Boolean duals.

Definition 3.36 Let $\varphi \in \text{FOE}_1^\infty(A)$. The *dual* $\varphi^\delta \in \text{FOE}_1^\infty(A)$ of φ is defined as follows.

$$\begin{aligned} (a(x))^\delta &:= a(x) & (\perp)^\delta &:= \top \\ (\top)^\delta &:= \perp & (x \not\approx y)^\delta &:= x \approx y \\ (x \approx y)^\delta &:= x \not\approx y & (\varphi \vee \psi)^\delta &:= (\varphi)^\delta \wedge (\psi)^\delta \\ (\varphi \wedge \psi)^\delta &:= (\varphi)^\delta \vee (\psi)^\delta & (\forall x.\psi)^\delta &:= \exists x.(\psi)^\delta \\ (\exists x.\psi)^\delta &:= \forall x.(\psi)^\delta & (\forall^\infty x.\psi)^\delta &:= \exists^\infty x.(\psi)^\delta \\ (\exists^\infty x.\psi)^\delta &:= \forall^\infty x.(\psi)^\delta \end{aligned}$$

Remark 3.37 Observe that if $\varphi \in \text{FO}_1(A)$ then $\varphi^\delta \in \text{FO}_1(A)$ and that the operator preserves positivity of the predicates. That is, if $\varphi \in \text{FOE}_1^{\infty+}(A)$ then $\varphi^\delta \in \text{FOE}_1^{\infty+}(A)$ and the same occurs with $\text{FO}_1^+(A)$. \triangleleft

The proof of the following Proposition is a routine check.

Proposition 3.38 *The sentences φ and φ^δ are Boolean duals, for every $\varphi \in \text{FOE}_1^\infty(A)$.*

We are now ready to give the syntactic definition of a co-continuous fragment for the one-step logics into consideration.

Definition 3.39 Let A be a set of names. The syntactic fragments of $\text{FOE}_1^\infty(A)$ and $\text{FO}_1(A)$ which are *co-continuous* in $a \in A$ are given by

$$\text{FOE}_1^\infty \overline{C}_a(A) := \{\varphi \mid \varphi^\delta \in \text{FOE}_1^\infty C_a(A)\} \quad \text{FO}_1 \overline{C}_a(A) := \{\varphi \mid \varphi^\delta \in \text{FO}_1 C_a(A)\}.$$

Proposition 3.40 *A formula $\varphi \in \text{FOE}_1^\infty(A)$ is co-continuous in $a \in A$ iff $\varphi \in \text{FOE}_1^\infty \overline{C}_a(A)$.*

Proof. This is a consequence of Proposition 3.35, Theorem 3.31 and Definition 3.39. QED

4 Automata for WMSO

In this section we provide an automata-theoretic characterization of WMSO. As argued in the introduction, WMSO-automata will be defined as the automata in $\text{Aut}(\text{FOE}_1^\infty)$ that satisfy two additional properties: weakness and continuity. We now briefly discuss these properties in a slightly more general setting, starting with weakness.

Definition 4.1 Let \mathcal{L}_1 be a one-step language, and let $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ be in $\text{Aut}(\mathcal{L}_1)$. Given two states a, b of \mathbb{A} , we say that there is a transition from a to b , notation: $a \rightsquigarrow b$, if b occurs in $\Delta(a, c)$ for some $c \in C$. We let the *reachability* relation \preceq denote the reflexive-transitive closure of the relation \rightsquigarrow .

A *strongly connected \preceq -component* (\preceq -SCC) is a subset $M \subseteq A$ such that for every $a, b \in M$ we have $a \preceq b$ and $b \preceq a$. The SCC is called *maximal* (MSCC) when $M \cup \{a\}$ ceases to be a SCC for any choice of $a \in A \setminus M$.

We say that Ω is a *weak parity condition*, and \mathbb{A} is a *weak parity automaton* if we have (**weakness**) if $a \preceq b$ and $b \preceq a$ then $\Omega(a) = \Omega(b)$.

Remark 4.2 Any weak parity automaton \mathbb{A} is equivalent to a weak parity automaton \mathbb{A}' with $\Omega : A' \rightarrow \{0, 1\}$. From now on we therefore consider only weak parity automata with priorities 0 and 1. \triangleleft

As explained in the introduction, the leading intuition is that weak parity automata are those unable to register non-trivial properties concerning the vertical ‘dimension’ of input trees. Indeed on trees they characterize WFMSO, that is, the fragment of MSO where the quantification is restricted to *well-founded* subsets of trees (corresponding to the notion of

noetherian subset if we consider arbitrary transition systems). We refer to the literature on weak automata [14] and on WFMSO [5, 24] for more details.

We now turn to the second condition that we will be interested in, viz., continuity. Intuitively, this property expresses a constraint on how much of the horizontal ‘dimension’ of an input tree the automaton is allowed to process. It will be instrumental in showing that the class of automata that we are going to shape characterizes WMSO. The idea is that, as WMSO-quantifiers range over finite sets, the weakness condition corresponds to those sets being ‘vertically’ finite (i.e. included in well-founded subtrees), whereas the continuity condition corresponds to them being ‘horizontally’ finite (i.e. included in finitely branching subtrees).

First we formulate our continuity condition abstractly in the setting of $Aut(\mathcal{L}_1)$. Given the semantics of the one-step language \mathcal{L}_1 , the (semantic) notion of (co-)continuity applies to one-step formulas (see for instance section 3.3). We can then formulate the following requirement on automata from $Aut(\mathcal{L}_1)$:

(continuity) let a, b be states such that both $a \preceq b$ and $b \preceq a$, and let $c \in C$; if $\Omega(b) = 1$ then $\Delta(b, c)$ is continuous in a . If $\Omega(b) = 0$, then $\Delta(b, c)$ is co-continuous in a .

For the automata used in this article we need to combine the constraints for the horizontal and vertical dimensions, yielding automata with both the weakness and continuity constraints.

Definition 4.3 A *continuous-weak parity automaton* is an automaton $\mathbb{A} \in Aut(\mathcal{L}_1)$ additionally satisfying both the **(weakness)** and **(continuity)** conditions. We let $Aut_{cw}(\mathcal{L}_1)$ denote the class of such automata.

Observe that, so far, the continuity condition is given semantically. However, given that the one-step languages that we are interested in have a *syntactic characterization* of continuity (see for example Theorem 3.31) we will give concrete definitions of these automata that take advantage of the mentioned characterizations.

Definition 4.4 A *WMSO-automaton* $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ is an automaton $\mathbb{A} \in Aut(\text{FOE}_1^\infty)$ such that for all states $a, b \in A$ with $a \preceq b$ and $b \preceq a$ the following conditions hold:

(weakness) $\Omega(a) = \Omega(b)$,

(continuity) if $\Omega(a)$ is odd (resp. even) then, for each $c \in C$ we have $\Delta(a, c) \in \text{FOE}_1^{\infty+} C_b(A)$ (resp. $\Delta(a, c) \in \text{FOE}_1^{\infty+} \overline{C}_b(A)$).

As the class of such automata coincides with $Aut_{cw}(\text{FOE}_1^\infty)$ we use the same notation for it.

We have now arrived at the main theorem of this section, which takes care of one direction of Theorem 1.2. The main theorem of this section states that

Theorem 4.5 *There is an effective construction transforming a WMSO-formula φ into a WMSO-automaton \mathbb{A}_φ that is equivalent to φ on the class of trees. That is, for any tree \mathbb{T} ,*

$$\mathbb{A}_\varphi \text{ accepts } \mathbb{T} \text{ iff } \mathbb{T} \models \varphi. \quad (10)$$

As usual, the proof of this theorem proceeds by induction on the complexity of φ . For the inductive steps of this proof, we need to verify that the class of WMSO-automata is closed under the boolean operations and finite projection. Clearly, the latter closure property requires most of the work; we first provide a simulation theorem that put WMSO-automata in a suitable shape for the projection construction. The proofs in the next two sections are (nontrivial) modifications of the analogous results proved in [5]. It will be convenient to use the following terminology.

Definition 4.6 The *tree language* $\mathcal{T}(\mathbb{A})$ recognized by an automaton \mathbb{A} is defined as the collection of trees that are accepted by \mathbb{A} . A class of trees is *WMSO-automaton recognizable* (or *recognizable*, if clear from context), if it is of the form $\mathcal{T}(\mathbb{A})$ for a WMSO-automaton \mathbb{A} .

4.1 Simulation theorem

In this section we show that any WMSO-automaton \mathbb{A} can be simulated by a “two-sorted” WMSO-automaton \mathbb{A}^F . The leading intuition is that \mathbb{A}^F will consist of one copy of \mathbb{A} (based on a set of states A) and a variant of its powerset construction, which will be based both on states from A and “macro-states” from $\wp(A \times A)$.³ Successful runs of \mathbb{A}^F will have the property of processing only a *finite* amount of the input with \mathbb{A}^F being in a macro-state and all the rest with \mathbb{A}^F behaving exactly as \mathbb{A} .

To achieve this result, we first need some preliminary definitions. The following is a notion of lifting for types on states that is instrumental in defining a translation to types on macro-states. The distinction between empty and non-empty subsets of A is to make sure that empty types on A are lifted to empty types on $\wp A$.

Definition 4.7 Let A be a set of unary predicates. Given a set $\Sigma \subseteq \wp A$, its *lifting* is the set $\Sigma^\wp \subseteq \wp \wp A$ defined as follows:

$$\Sigma^\wp := \{\{S\} \mid S \in \Sigma \wedge S \neq \emptyset\} \cup \{\emptyset \mid \emptyset \in \Sigma\}.$$

The next step is to define a translation on the sentences associated with the transition function of the original WMSO-automaton, say with set of states A . Following the intuition given above, the idea is that we want to work with sentences that can be made true by assigning macro-states (from $\wp(A \times A)$) to finitely many nodes in the model, and ordinary states (from A) to all the other nodes. Henceforth, we use the notation A^\sharp for the set $\wp(A \times A)$.

Definition 4.8 Let $\varphi \in \text{FOE}_1^{\infty+}(A \times A)$ be a formula of shape $\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$ for some $\Pi, \Sigma \subseteq A^\sharp$ and $\vec{T} = \{T_1, \dots, T_k\} \subseteq A^\sharp$. Let $\tilde{\Sigma} \subseteq \wp A$ be $\tilde{\Sigma} := \{\text{Ran}(S) \mid S \in \Sigma\}$. We define $\varphi^F \in \text{FOE}_1^{\infty+}(A \cup A^\sharp)$ as follows:

$$(\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma))^F := \nabla_{\text{FOE}^\infty}^+(\vec{T}^\wp, \Pi^\wp \cup \Sigma^\wp, \tilde{\Sigma}).$$

³It is customary for powerset constructions on parity automata to encode macro-states as binary relations between states (from $\wp(A \times A)$) instead of plain sets (from $\wp A$). Such additional structure is needed to correctly associate with a run on macro-states the corresponding bundle of runs of the original automaton \mathbb{A} . We refer to the standard literature on parity automata (e.g. [23, 8]) for further details.

Observe that each τ_P^+ with $P \in \tilde{\Sigma}$ appearing in φ^F is a (positive) A -type, as $P = \text{Ran}(S) \subseteq A$ for some $S \in \Sigma$.

Our desiderata on the translation $(\cdot)^F$ concern the notions of *continuity* and *functionality*.

Definition 4.9 Given a set A of unary predicates and $B \subseteq A$, we say that a sentence $\varphi \in \text{FOE}_1^{\infty+}(A)$ is *functionally continuous in B* if, for every model $(D, V: A \rightarrow \wp(D))$,

- if $(D, V), g \models \varphi$ then $\exists V': A \rightarrow \wp(D)$ such that $(D, V'), g \models \varphi$,
- $V'(a) \subseteq V(a)$ for all $a \in A$, (V' is a restriction of V)
 - $V'(b)$ is finite for all $b \in B$ and (continuity in B)
 - $V'(b) \cap V'(a) = \emptyset$ for all $a \in A \setminus \{b\}$ and $b \in B$. (functionality in B)

In words, φ is functionally continuous in B if it is continuous in each $b \in B$ and, for each model (D, V) where φ is true, there is a restriction V' of V which both witnesses continuity and does not assign any other $a \in A$ to the elements marked with some $b \in B$.

Lemma 4.10 Let $\varphi \in \text{FOE}_1^{\infty+}(A \times A)$ and $\varphi^F \in \text{FOE}_1^{\infty+}(A \cup A^\#)$ be given as in Definition 4.8. Then φ^F is functionally continuous in $A^\#$.

Proof. We first unfold the definition of φ^F as follows:

$$\begin{aligned} \varphi^F = & \underbrace{\exists \vec{x}. (\text{diff}(\vec{x}) \wedge \bigwedge_{0 \leq i \leq n} \tau_{T_i^\wp}^+(x_i) \wedge \forall z. (\text{diff}(\vec{x}, z) \rightarrow \bigvee_{S \in \Pi^\wp \cup \Sigma^\wp \cup \tilde{\Sigma}} \tau_S^+(z)))}_{\psi_1} \wedge \underbrace{\bigvee_{S \in \Pi^\wp \cup \Sigma^\wp \cup \tilde{\Sigma}} \tau_S^+(z)}_{\psi_2} \\ & \underbrace{\bigwedge_{P \in \tilde{\Sigma}} \exists^\infty y. \tau_P^+(y)}_{\psi_3} \wedge \underbrace{\forall^\infty y. \bigvee_{P \in \tilde{\Sigma}} \tau_P^+(y)}_{\psi_4}. \end{aligned}$$

Observe that $\psi_1 \wedge \psi_2$ is just $\nabla_{\text{FOE}}^+(\vec{T}^\wp, \Pi^\wp \cup \Sigma^\wp \cup \tilde{\Sigma})$. Now suppose that $(D, V: (A \cup A^\#) \rightarrow \wp(D))$ is a model where φ^F is true. This amounts to the truth of subformulas ψ_1, ψ_2, ψ_3 and ψ_4 whose syntactic shape yields information on the types of elements of D . In particular, we can define a partition of D into subsets D_1, D_2, D'_2 as follows:

- As ψ_1 is true, we can pick n distinct elements s_1, \dots, s_n of D such that s_i witnesses the positive type T_i^\wp , that is, $s_i \in V(S)$ for each $S \in T_i^\wp$. We define $D_1 := \{s_1, \dots, s_n\}$.
- As ψ_2 is true, we can cover all the elements not in D_1 with two disjoint sets D_2 and D'_2 given as follows. The set D_2 is defined to contain all the elements not in D_1 witnessing a type $\tau_P^+(z)$ with $P \in \tilde{\Sigma}$. The set D'_2 is just the complement of $D_1 \cup D_2$: by syntactic shape of ψ_2 , all elements of D'_2 witness a positive type τ_S^+ with $S \in \Pi^\wp \cup \Sigma^\wp$.
- The truth of the subformula ψ_4 yields the information that the set $D_1 \cup D'_2$ is finite. If $\tilde{\Sigma}$ is non-empty, the truth of ψ_3 implies that the set D_2 is infinite.

This partition uniquely associates with each $s \in D$ a type τ_S^+ witnessed by s and thus a set of unary predicates $S_s := S \subseteq A \cup A^\sharp$. We can then define a valuation V' assigning to each element s of D exactly the set S_s .

We now check the properties of V' . As the partition inducing V' follows the syntactic shape of φ^F , one can observe that V' is a restriction of V and (D, V') makes φ^F true. By definition of the partition, V' assigns unary predicates from A^\sharp only to elements in the finite set $D_1 \cup D'_2$, meaning that φ^F is continuous in A^\sharp . Furthermore, V' assigns at most one unary predicate from A^\sharp to each element of $D_1 \cup D'_2$, because $\vec{T}^\wp \cup \Pi^\wp \cup \Sigma^\wp$ is defined as the lifting of $\vec{T} \cup \Pi \cup \Sigma$. It follows that φ^F is also functional in A^\sharp . Since the same restriction V' yields both properties, φ^F is functionally continuous in A^\sharp . QED

Remark 4.11 As φ^F is of shape $\nabla_{\text{FOE}^\infty}^+(\vec{T}^\wp, \Pi^\wp \cup \Sigma^\wp, \tilde{\Sigma})$ with $R \notin \bigcup \tilde{\Sigma}$ for each $R \in A^\sharp$, by application of Corollary 3.34 we would immediately get that φ^F is continuous in each $R \in A^\sharp$. However, we do not use this observation in proving Lemma 4.10 and propose instead a more direct argument, allowing to show both continuity and functionality at once. \triangleleft

The next definition is standard (see e.g. [23, 22]) as an intermediate step to define the transition function of the powerset construct for parity automata.

Definition 4.12 Let $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ be a WMSO-automaton. Fix $a \in A$ and $c \in C$. The sentence $\Delta^*(a, c)$ is defined as

$$\Delta^*(a, c) \quad := \quad \Delta(a, c)[b \mapsto (a, b) \mid b \in A].$$

Next we combine the previous definitions to characterize the transition function associated with the macro-states.

Definition 4.13 Let $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ be a WMSO-automaton. Let $c \in C$ be a label and $Q \in A^\sharp$ a binary relation on A . By Corollary 3.26, for some $\Pi, \Sigma \subseteq A^\sharp$ and $T_i \subseteq A \times A$, there is a sentence $\Psi_{Q,c} \in \text{FOE}_1^{\infty+}(A \times A)$ in the basic form $\bigvee \nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$ such that

$$\bigwedge_{a \in \text{Ran}(Q)} \Delta^*(a, c) \quad \equiv \quad \Psi_{Q,c}.$$

By definition $\Psi_{Q,c}$ is of the form $\bigvee_i \varphi_i$, with each φ_i of shape $\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$. We put $\Delta^\sharp(Q, c) := \bigvee_i \varphi_i^F$, where the translation $(\cdot)^F$ is given as in definition 4.8. Observe that $\Delta^\sharp(Q, c)$ is of type $\text{FOE}_1^{\infty+}(A \cup A^\sharp)$.

We have now all the ingredients to define our two-sorted automaton.

Definition 4.14 Let $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ be a WMSO-automaton. We define the *finitary construct over \mathbb{A}* as the automaton $\mathbb{A}^F = \langle A^F, \Delta^F, \Omega^F, a_I^F \rangle$ given by

$$\begin{aligned} A^F &:= A \cup A^\sharp \\ a_I^F &:= \{(a_I, a_I)\} \\ \Delta^F(a, c) &:= \Delta(a, c) \\ \Delta^F(R, c) &:= \Delta^\sharp(R, c) \vee \bigwedge_{a \in \text{Ran}(R)} \Delta(a, c) \\ \Omega^F(a) &:= \Omega(a) \\ \Omega^F(R) &:= 1. \end{aligned}$$

The underlying idea of Definition 4.14 is the same of the *two-sorted construction* (cf. [24, Def. 3.7], [5]) for weak MSO-automata. In both cases we want that macro-states process just a *well-founded* portion of any accepted tree: this is guaranteed by associating all macro-states with the odd parity value 1. However, for the finitary construction we aim at the stronger condition that such a portion is *finite*. To achieve this, the key difference with the two-sorted construction is in the use of the translation $(\cdot)^F$ to define Δ^\sharp : as Δ may be specified using quantifiers \exists^∞ and \forall^∞ , it serves the purpose of tracking the cardinality constraints of the original WMSO-automaton and ensure that they are not lifted to constraints on macro-states.

The next proposition establishes the desired properties of the finitary construct. To this aim, we first introduce the notions of functional and finitary strategy.

Definition 4.15 Given a WMSO-automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ and transition system \mathbb{T} , a strategy f for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{T})$ is *functional in $B \subseteq A$* (or simply functional, if $B = A$) if for each node s in \mathbb{T} there is at most one $b \in B$ such that (b, s) is a reachable position in an f -guided match. Also f is *finitary in B* if there are only finitely many nodes s in \mathbb{T} for which a position (b, s) with $b \in B$ is reachable in an f -guided match.

Lemma 4.16 Let \mathbb{A} be a WMSO-automaton and \mathbb{A}^F its finitary construct over \mathbb{A} . The following holds:

1. \mathbb{A}^F is a WMSO-automaton.
2. For any \mathbb{T} , if \exists has a winning strategy in the game $\mathcal{A}(\mathbb{A}^F, \mathbb{T}) @ (a_I^F, s_I)$, then she has a winning strategy in the same game which is both functional and finitary in A^\sharp .
3. $\mathbb{A} \equiv \mathbb{A}^F$.

Proof. We address each point separately.

1. We need to show that \mathbb{A}^F is weak and respects the continuity condition. For this purpose, we fix the following observation:

(\star) by definition of Δ^F , for any macro-state $R \in A^\sharp$ and state $a \in A$, it is never the case that $a \preceq R$.

This means that, when considering a strongly connected component of \mathbb{A}^F , we may assume that all states involved are either from A or from A^\sharp .

In order to prove our claim, let $q_1, q_2 \in A^F$ be two states of \mathbb{A}^F such that $q_1 \preceq q_2$ and $q_2 \preceq q_1$. By observation (\star) , we can distinguish the following two cases:

- (i) if q_1 and q_2 are states from A , then for $i \in \{1, 2\}$ the value of $\Omega^F(q_i)$ and $\Delta^F(q_i, c)$ is defined respectively as $\Omega(q_i)$ and $\Delta(q_i, c)$ like in the WMSO-automaton \mathbb{A} . It follows that they satisfy both the continuity and weakness condition.
- (ii) Otherwise, q_1 and q_2 are macro-states in A^\sharp . For the weakness condition, observe that all macro-states in \mathbb{A}^F have the same parity value. For the continuity condition, suppose that q_1 occurs in $\Delta^F(q_2, c)$ for some $c \in C$. By definition q_1 can only appear in the disjunct $\Delta^\sharp(q_2, c) = \bigvee_i \varphi_i^F$ of $\Delta^F(q_2, c)$. By Lemma 4.10, we know that each φ_i^F is continuous in A^\sharp . Then in particular $\Delta^F(q_2, c)$ is continuous in q_1 . By definition $\Omega^F(q_1) = 1$ is odd, meaning that the continuity condition holds. The case in which q_2 appears in $\Delta^F(q_1, c)$ is just symmetric.

2. Let f be a winning strategy for \exists in $\mathcal{A}(\mathbb{A}^F, \mathbb{T}) @ (a_I^F, s_I)$. We define a strategy f' for \exists in the same game as follows:

- (a) on basic positions of the form $(a, s) \in A \times T$, let V be the valuation suggested by f . We let the valuation suggested by f' be the restriction V' of V to A . Observe that, as no predicate from $A^F \setminus A = A^\sharp$ occurs in $\Delta^F(a, \sigma(s)) = \Delta(a, \sigma(s))$, then V' also makes that sentence true in $R[s]$.
- (b) for basic positions of the form $(R, s) \in A^\sharp \times T$, let $V_{R,s}$ be the valuation suggested by f . As f is winning, $\Delta^F(R, \sigma(s))$ is true in the model $V_{R,s}$. If this is because the disjunct $\bigwedge_{a \in \text{Ran}(R)} \Delta(a, \sigma(s))$ is made true, then we can let f' suggest the restriction to A of $V_{R,s}$, for the same reason as in (a). Otherwise, the disjunct $\Delta^\sharp(R, \sigma(s)) = \bigvee_i \varphi_i^F$ is made true. This means that, for some i ,

$$(R[s], V_{R,s}) \models \varphi_i^F.$$

By Lemma 4.10 φ_i^F is functionally continuous in A^\sharp , meaning that we have a restriction $V'_{R,s}$ of $V_{R,s}$ that verifies φ_i^F , assigns finitely many nodes to predicates from A^\sharp and associates with each node at most one predicate from A^\sharp . We let $V'_{R,s}$ be the suggestion of f' from position (R, s) .

The strategy f' defined as above is immediately seen to be surviving for \exists . It is also winning, because the set of basic positions on which f' is defined is a subset of the one of the winning strategy f . By this observation it also follows that any f' -conform match visits basic positions of the form $(R, s) \in A^\sharp \times C$ only finitely many times, as those have odd parity. By definition, the valuation suggested by f' only assigns finitely many nodes to predicates in A^\sharp from positions of that shape, and no nodes from other positions. It follows that f' is finitary in A^\sharp . Functionality in A^\sharp also follows immediately by definition of f' .

3. The proof is entirely analogous to the one presented in [24, Prop. 3.9] for the two-sorted construction. For the direction from left to right, it is immediate by definition of \mathbb{A}^F that a winning strategy for \exists in $\mathcal{G} = \mathcal{A}(\mathbb{A}, \mathbb{T}) @ (a_I, s_I)$ is also winning for \exists in $\mathcal{G}^F = \mathcal{A}(\mathbb{A}^F, \mathbb{T}) @ (a_I^F, s_I)$.

For the direction from right to left, let f be a winning strategy for \exists in \mathcal{G}^F . The idea is to define a strategy f' for \exists in stages, while playing a match π' in \mathcal{G} . In parallel to π' , a shadow match π in \mathcal{G}^F is maintained, where \exists plays according to the strategy f . For each round z_i , we want to keep the following relation between the two matches:

<p>Either</p> <p>(1) basic positions of the form $(Q, s) \in A^\sharp \times T$ and $(a, s) \in A \times T$ occur respectively in π and π', with $a \in \text{Ran}(Q)$,</p> <p>or</p> <p>(2) the same basic position of the form $(a, s) \in A \times T$ occurs in both matches.</p>	(\ddagger)
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------

The key observation is that, because f is winning, a basic position of the form $(Q, s) \in A^\sharp \times T$ can occur only for finitely many initial rounds z_0, \dots, z_n that are played in π , whereas for all successive rounds z_n, z_{n+1}, \dots only basic positions of the form $(a, s) \in A \times T$ are encountered. Indeed, if this was not the case then either \exists would get stuck or the minimum parity occurring infinitely often would be odd, since states from A^\sharp have parity 1.

It follows that enforcing a relation between the two matches as in (\ddagger) suffices to prove that the defined strategy f' is winning for \exists in π' . For this purpose, first observe that $(\ddagger).1$ holds at the initial round, where the positions visited in π' and π are respectively $(a_I, s_I) \in A \times T$ and $(\{(a_I, a_I)\}, s_I) \in A^F \times T$. Inductively, consider any round z_i that is played in π' and π , respectively with basic positions $(a, s) \in A \times T$ and $(q, s) \in A^F \times T$. In order to define the suggestion of f' in π' , we distinguish two cases.

- First suppose that (q, s) is of the form $(Q, s) \in A^\sharp \times T$. By (\ddagger) we can assume that a is in $\text{Ran}(Q)$. Let $V_{Q,s} : A^F \rightarrow \wp(R[s])$ be the valuation suggested by f , verifying the sentence $\Delta^F(Q, \sigma(s))$. We distinguish two further cases, depending on which disjunct of $\Delta^F(Q, \sigma(s))$ is made true by $V_{Q,s}$.
 - (i) If $(R[s], V_{Q,s}) \models \bigwedge_{b \in \text{Ran}(Q)} \Delta(b, \sigma(s))$, then we let \exists pick the restriction to A of the valuation $V_{Q,s}$.
 - (ii) If $(R[s], V_{Q,s}) \models \Delta^\sharp(Q, \sigma(s))$, we let \exists pick a valuation $V_{a,s} : A \rightarrow \wp(R[s])$ defined by putting, for each $b \in A$:

$$V_{a,s}(b) := \bigcup_{b \in \text{Ran}(Q')} \{t \in R[s] \mid t \in V_{Q,s}(Q')\} \cup \{t \in R[s] \mid t \in V_{Q,s}(b)\}.$$

It can be readily checked that the suggested move is admissible for \exists in π , i.e. it makes $\Delta(a, \sigma(s))$ true in $R[s]$. For case (ii), one has to observe how Δ^\sharp is defined in terms of Δ . In particular, the nodes assigned to b by $V_{Q,s}$ have to be assigned to b also by $V_{a,s}$, as they may be necessary to fulfill the condition, expressed with \exists^∞ and \forall^∞ , that infinitely many nodes witness (or that finitely many nodes do not witness) some type.

We now show that (\ddagger) holds at round z_{i+1} . If (i) is the case, any next position $(b, t) \in A \times T$ picked by player \forall in π' is also available for \forall in π , and we end up in case $(\ddagger.2)$. Suppose instead that (ii) is the case. Given the choice $(b, t) \in A \times T$ of \forall , by definition of $V_{a,s}$ there are two possibilities. First, (b, t) is also an available choice for \forall in π , and we end up in case $(\ddagger.2)$ as before. Otherwise, there is some $Q' \in A^\sharp$ such that b is in $\text{Ran}(Q')$ and \forall can choose (Q', t) in the shadow match π . By letting π advance at round z_{i+1} with such a move, we are able to maintain $(\ddagger.1)$ also in z_{i+1} .

- In the remaining case, inductively we are given the same basic position $(a, s) \in A \times T$ both in π and in π' . The valuation V suggested by f in π verifies $\Delta^F(a, \sigma(s)) = \Delta(a, \sigma(s))$, thus we can let the restriction of V to A be the valuation chosen by \exists in the match π' . It is immediate that any next move of \forall in π' can be mirrored by the same move in π , meaning that we are able to maintain the same position –whence the relation $(\ddagger.1)$ – also in the next round.

In both cases, the suggestion of strategy f' was a legitimate move for \exists maintaining the relation (\ddagger) between the two matches for any next round z_{i+1} . It follows that f' is a winning strategy for \exists in \mathcal{G} .

QED

Remark 4.17 While the finitary construction is a variant of the two-sorted one given in [5], it is worth noticing that the latter would have not been suitable for our purposes. Indeed, suppose to define the two-sorted construct \mathbb{A}^{2S} over a WMSO-automaton \mathbb{A} , analogously to the case of weak MSO-automata. Then \mathbb{A}^{2S} will generally *not* be a WMSO-automaton. The problem lies in the **(continuity)** condition: since all macro-states in \mathbb{A}^{2S} have parity 1, whenever two of them, say R and Q , are such that $R \preceq Q$ and $Q \preceq R$, then the sentence $\Delta^{2S}(R, c)$ should be continuous in Q . But this is not necessarily the case, since the truth of $\Delta^{2S}(R, c)$ may depend upon the truth of a subformula of the form $\exists^\infty x.Q(x)$, requiring Q to be interpreted over infinitely many nodes. (This problem is overcome in the finitary construction by using the translation $(\cdot)^F$ to define Δ^F .)

As a consequence, we cannot use the two-sorted construction to show that WMSO-automata are closed under noetherian projection ([5, Def. 3]). This observation is coherent with the fact that WFMSO is *not* a fragment of WMSO. Similarly, the simulation theorem for MSO-automata [23] preserves neither the **(weakness)** nor the **(continuity)** condition and thus it cannot show closure under (arbitrary) projection for WMSO-automata. \triangleleft

4.2 Closure Properties

In this subsection we prove that WMSO-automata are closed under the operations corresponding to the connectives of MSO, that is: union, complementation and projection with respect to finite sets. We start with the latter.

4.2.1 Closure under Finitary Projection

Definition 4.18 Let $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ be a WMSO-automaton on alphabet $\wp(\mathbf{P} \cup \{p\})$. Let \mathbb{A}^F denote its finitary construct. We define the automaton $\exists_{Fp}.\mathbb{A} = \langle A^F, a_I^F, \tilde{\Delta}, \Omega^F \rangle$ on alphabet $\wp\mathbf{P}$ by putting

$$\begin{aligned}\tilde{\Delta}(a, c) &:= \Delta^F(a, c) \\ \tilde{\Delta}(R, c) &:= \Delta^F(R, c) \vee \Delta^F(R, c \cup \{p\}).\end{aligned}$$

The automaton $\exists_{Fp}.\mathbb{A}$ is called the *finitary projection construct* of \mathbb{A} over p .

Our projection construction corresponds to a suitable closure operation on tree languages, modeling the semantics of WMSO existential quantification.

Definition 4.19 Let p be a propositional letter and L a tree language of $\wp(\mathbf{P} \cup \{p\})$ -labeled trees. The *finite projection* of L over p is the language $\exists_{Fp}.L$ of C -labeled trees defined as

$$\exists_{Fp}.L = \{\mathbb{T} \mid \text{there is a } p\text{-variant } \mathbb{T}[p \mapsto S] \text{ of } \mathbb{T} \text{ such that } \mathbb{T}[p \mapsto S] \in L \text{ and } S \text{ is finite}\}.$$

Lemma 4.20 For each WMSO-automaton \mathbb{A} on alphabet $\wp(\mathbf{P} \cup \{p\})$, we have that

$$\mathcal{T}(\exists_{Fp}.\mathbb{A}) \equiv \exists_{Fp}.\mathcal{T}(\mathbb{A}).$$

Proof. What we need to show is that for any tree \mathbb{T} :

$$\begin{aligned}\exists_{Fp}.\mathbb{A} \text{ accepts } \mathbb{T} \quad \text{iff} \quad & \text{there is a finite } p\text{-variant } \mathbb{T}' \\ & \text{of } \mathbb{T} \text{ such that } \mathbb{A} \text{ accepts } \mathbb{T}'.\end{aligned}$$

For direction from left to right, we first observe that the properties stated by Lemma 4.16 hold for $\exists_{Fp}.\mathbb{A}$ as well, since the latter is defined in terms of \mathbb{A}^F . Then we can assume that the given winning strategy f for \exists in $\mathcal{G}_{\exists} = \mathcal{A}(\exists_{Fp}.\mathbb{A}, \mathbb{T}) @ (a_I^F, s_I)$ is functional and finitary in A^\sharp . Functionality allows us to associate with each node s either none or a unique state $Q_s \in A^\sharp$ (cf. [24, Prop. 3.12]). We now want to isolate the nodes that f treats “as if they were labeled with p ”. For this purpose, let V_s be the valuation suggested by f from a position $(Q_s, s) \in A^\sharp \times T$. As f is winning, V_s makes $\tilde{\Delta}(Q, \sigma(s))$ true in $R[s]$. We define a p -variant \mathbb{T}' of \mathbb{T} by coloring with p all nodes in the following set:

$$X_p := \{s \in T \mid (R[s], \tilde{V}_s) \models \Delta^F(Q_s, \sigma(s) \cup \{p\})\}. \quad (11)$$

The fact that the strategy of \exists is finitary in A^\sharp guarantees that X_p is finite, whence \mathbb{T}' is a finite p -variant. The argument showing that \mathbb{A}^F (and thus also \mathbb{A} , by Lemma 4.16(1))

accepts \mathbb{T}' is a routine adaptation of the analogous proof for the noetherian projection of weak MSO-automata, for which we refer to [24, Prop. 3.12].

For the direction from right to left, let \mathbb{T}' be a finite p -variant of \mathbb{T} , with labeling function σ' , and g a winning strategy for \exists in $\mathcal{G} = \mathcal{A}(\mathbb{A}, \mathbb{T}') @ (a_I, s_I)$. Our goal is to define a strategy g' for \exists in \mathcal{G}_\exists . As usual, g' will be constructed in stages, while playing a match π' in \mathcal{G}_\exists . In parallel to π' , a *bundle* \mathcal{B} of g -guided shadow matches in \mathcal{G} is maintained, with the following condition enforced for each round z_i (cf. [24, Prop. 3.12]) :

1. If the current (i.e. at round z_i) basic position in π' is of the form $(Q, s) \in A^\sharp \times T$, then for each $a \in \text{Ran}(Q)$ there is an g -guided (partial) shadow match π_a at basic position $(a, s) \in A \times T$ in the current bundle \mathcal{B}_i . Also, either \mathbb{T}'_s is not p -free (i.e., it does contain a node s' with $p \in \sigma'(s')$) or s has some sibling t such that \mathbb{T}'_t is not p -free.
 2. Otherwise, the current basic position in π' is of the form $(a, s) \in A \times T$ and \mathbb{T}'_s is p -free (i.e., it does not contain any node s' with $p \in \sigma'(s')$). Also, the bundle \mathcal{B}_i only consists of a single g -guided match π_a whose current basic position is also (a, s) .
- (‡)

We briefly recall the idea behind condition (‡). Point (‡.1) describes the part of match π' where it is still possible to encounter nodes which are labeled with p in \mathbb{T}' . As $\tilde{\Delta}$ only takes the letter p into account when defined on macro-states in A^\sharp , we want π' to visit only positions of the form $(R, s) \in A^\sharp \times T$ in that situation. Anytime we visit such a position (R, s) in π' , the role of the bundle is to provide one g -guided shadow match at position (a, s) for each $a \in \text{Ran}(R)$. Then g' is defined in terms of what g suggests from those positions.

Point (‡.2) describes how we want the match π' to be played on a p -free subtree: as any node that one might encounter has the same label in \mathbb{T} and \mathbb{T}' , it is safe to let $\exists_F p. \mathbb{A}$ behave as \mathbb{A} in such situation. Provided that the two matches visit the same basic positions, of the form $(a, s) \times A \times T$, we can let g' just copy g .

The key observation is that, as \mathbb{T}' is a *finite* p -variant of \mathbb{T} , nodes labeled with p are reachable only for finitely many rounds of π' . This means that, provided that (‡) hold at each round, (‡.1) will describe an initial segment of π' , whereas (‡.2) will describe the remaining part. Thus our proof that g' is a winning strategy for \exists in \mathcal{G}_\exists is concluded by showing that (‡) holds for each stage of construction of π' and \mathcal{B} .

For this purpose, we initialize π' from position $(a_I^\sharp, s) \in A^\sharp \times T$ and the bundle \mathcal{B} as $\mathcal{B}_0 = \{\pi_{a_I}\}$, with π_{a_I} the partial g -guided match consisting only of the position $(a_I, s) \in A \times T$. The situation described by (‡.1) holds at the initial stage of the construction. Inductively, suppose that at round z_i we are given a position $(q, s) \in A^F \times T$ in π^F and a bundle \mathcal{B}_i as in (‡). To show that (‡) can be maintained at round z_{i+1} , we distinguish two cases, corresponding respectively to situation (‡.1) and (‡.2) holding at round z_i .

- (A) If (q, s) is of the form $(Q, s) \in A^\sharp \times T$, by inductive hypothesis we are given with g -guided shadow matches $\{\pi_a\}_{a \in \text{Ran}(Q)}$ in \mathcal{B}_i . For each match π_a in the bundle, we are

provided with a valuation $V_{a,s} : A \rightarrow \wp(R[s])$ making $\Delta(a, \sigma'(s))$ true. Then we further distinguish the following two cases.

- (i) Suppose first that \mathbb{T}'_s is not p -free. We let the suggestion $V' : A^F \rightarrow \wp(R[s])$ of g' from position (Q, s) be defined as follows:

$$V'(q') := \begin{cases} \bigcap_{\substack{(a,b) \in q', \\ a \in \text{Ran}(Q)}} \{t \in R[s] \mid t \in V_{a,s}(b)\} & q' \in A^\sharp \\ \bigcup_{a \in \text{Ran}(Q)} \{t \in R[s] \mid t \in V_{a,s}(q') \text{ and } \mathbb{T}' \cdot t \text{ is } p\text{-free}\} & q' \in A. \end{cases}$$

The definition of V' on $q' \in A^\sharp$ is standard (cf. [24, Prop. 2.21]) and guarantees a correspondence between the states assigned by the markings $\{V_{a,s}\}_{a \in \text{Ran}(Q)}$ and the macro-states assigned by V' . The definition of V' on $q' \in A$ aims at fulfilling the conditions, expressed via \exists^∞ and \forall^∞ , on the number of nodes in $R[s]$ witnessing (or not) some A -types. Those conditions are the ones that $\Delta^\sharp(Q, \sigma'(s))$ –and thus also $\Delta^F(Q, \sigma'(s))$ – “inherits” by $\bigwedge_{a \in \text{Ran}(R)} \Delta(a, \sigma'(s))$, by definition of Δ^\sharp . Notice that we restrict $V'(q')$ to the nodes $t \in V_{a,s}(q')$ such that $\mathbb{T}' \cdot t$ is p -free. As \mathbb{T}' is a *finite* p -variant, only *finitely many* nodes in $V_{a,s}(q')$ will not have this property. Therefore their exclusion, which is crucial for maintaining condition (\ddagger) (cf. case (a) below), does not influence the fulfilling of the cardinality conditions expressed via \exists^∞ and \forall^∞ in $\Delta^\sharp(Q, \sigma'(s))$.

On the base of these observations, one can check that V' makes $\Delta^\sharp(Q, \sigma'(s))$ –and thus also $\Delta^F(Q, \sigma'(s))$ –true in $R[s]$. In fact, to be a legitimate move for \exists in π' , V' should make $\tilde{\Delta}(Q, \sigma(s))$ true: this is the case, for $\Delta^F(Q, \sigma'(s))$ is either equal to $\Delta^F(Q, \sigma(s))$, if $p \notin \sigma'(s)$, or to $\Delta^F(Q, \sigma(s) \cup \{p\})$ otherwise. In order to check that we can maintain (\ddagger) , let $(q', t) \in A^F \times T$ be any next position picked by \forall in π' at round z_{i+1} . As before, we distinguish two cases:

- (a) If q' is in A , then, by definition of V' , \forall can choose (q', t) in some shadow match π_a in the bundle \mathcal{B}_i . We dismiss the bundle –i.e. make it a singleton– and bring only π_a to the next round in the same position (q', t) . Observe that, by definition of V' , $\mathbb{T}' \cdot t$ is p -free and thus $(\ddagger.2)$ holds at round z_{i+1} .
 - (b) Otherwise, q' is in A^\sharp . The new bundle \mathcal{B}_{i+1} is given in terms of the bundle \mathcal{B}_i : for each $\pi_a \in \mathcal{B}_i$ with $a \in \text{Ran}(Q)$, we look if for some $b \in \text{Ran}(q')$ the position (b, t) is a legitimate move for \forall at round z_{i+1} ; if so, then we bring π_a to round z_{i+1} at position (b, t) and put the resulting (partial) shadow match π_b in \mathcal{B}_{i+1} . Observe that, if \forall is able to pick such position (q', t) in π' , then by definition of V' the new bundle \mathcal{B}_{i+1} is non-empty and consists of an g -guided (partial) shadow match π_b for each $b \in \text{Ran}(q')$. In this way we are able to keep condition $(\ddagger.1)$ at round z_{i+1} .
- (ii) Let us now consider the case in which \mathbb{T}'_s is p -free. We let g' suggest the valuation V' that assigns to each node $t \in R[s]$ all states in $\bigcup_{a \in \text{Ran}(Q)} \{b \in A \mid t \in V_{a,s}(b)\}$. It can be checked that V' makes $\bigwedge_{a \in \text{Ran}(Q)} \Delta(a, \sigma'(s))$ – and then also $\Delta^F(Q, \sigma'(s))$

– true in $R[s]$. As $p \notin \sigma(s) = \sigma'(s)$, it follows that V' also makes $\tilde{\Delta}(Q, \sigma(s))$ true, whence it is a legitimate choice for \exists in π' . Any next basic position picked by \forall in π' is of the form $(b, t) \in A \times T$, and thus condition $(\dagger.2)$ holds at round z_{i+1} as shown in (i.a).

(B) In the remaining case, (q, s) is of the form $(a, s) \in A \times T$ and by inductive hypothesis we are given with a bundle \mathcal{B}_i consisting of a single f -guided (partial) shadow match π_a at the same position (a, s) . Let $V_{a,s}$ be the suggestion of \exists from position (a, s) in π_a . Since by assumption s is p -free, we have that $\sigma'(s) = \sigma(s)$, meaning that $\tilde{\Delta}(a, \sigma(s))$ is just $\Delta(a, \sigma(s)) = \Delta(a, \sigma'(s))$. Thus the restriction V' of V to A makes $\Delta(a, \sigma'(t))$ true and we let it be the choice for \exists in $\tilde{\pi}$. It follows that any next move made by \forall in $\tilde{\pi}$ can be mirrored by \forall in the shadow match π_a .

QED

4.2.2 Closure under Boolean operations

In this section we will show that the class of WMSO-automaton recognizable tree languages is closed under the Boolean operations. Start with closure under union, we just mention the following result, without providing the (completely routine) proof.

Theorem 4.21 *Let \mathbb{A}_0 and \mathbb{A}_1 be WMSO-automata. Then there is a WMSO-automaton \mathbb{A} such that $\mathcal{T}(\mathbb{A})$ is the union of $\mathcal{T}(\mathbb{A}_0)$ and $\mathcal{T}(\mathbb{A}_1)$.*

In order to prove closure under complementation, we crucially use that the one-step language FOE_1^∞ is closed under Boolean duals (cf. Proposition 3.38).

Theorem 4.22 *Let \mathbb{A} be an WMSO-automaton. Then the automaton $\overline{\mathbb{A}}$ defined in Definition 2.6 is a WMSO-automaton recognizing the complement of $\mathcal{T}(\mathbb{A})$.*

Proof. Since we already know that $\overline{\mathbb{A}}$ accepts exactly the transition systems that are rejected by \mathbb{A} , we only need to check that $\overline{\mathbb{A}}$ is indeed a WMSO-automaton. But this is straightforward: for instance, continuity can be checked by observing the self-dual nature of this property. QED

4.3 Proof of Theorem 4.5

Proof. The proof is by induction on φ .

- For the base case $\varphi = p \sqsubseteq q$, the corresponding WMSO-automaton is provided in [24, Ex. 2.6]. For the base case $\varphi = R(p, q)$, we give the corresponding WMSO-automaton

$\mathbb{A}_{R(p,q)} = \langle A, \Delta, \Omega, a_I \rangle$ below:

$$\begin{aligned}
A &:= \{a_0, a_1\} \\
a_I &:= a_0 \\
\Delta(a_0, c) &:= \begin{cases} \exists x. a_1(x) \wedge \forall y. a_0(y) & \text{if } p \in c \\ \forall x (a_0(x)) & \text{otherwise} \end{cases} \\
\Delta(a_1, c) &:= \begin{cases} \top & \text{if } q \in c \\ \perp & \text{otherwise} \end{cases} \\
\Omega(a_0) &:= 0 \\
\Omega(a_1) &:= 1.
\end{aligned}$$

Note that the MSO-automaton for $R(p, q)$ provided in [24, Ex. 2.5] is *not* a WMSO-automaton, as the continuity property does not hold.

- For the Boolean cases, where $\varphi = \psi_1 \vee \psi_2$ or $\varphi = \neg\psi$ we refer to the closure properties of recognizable tree languages, see Theorem 4.21 and Theorem 4.22, respectively.
- For the case $\varphi = \exists p.\psi$, consider the following chain of equivalences, where \mathbb{A}_ψ is given by the inductive hypothesis and $\exists_F p.\mathbb{A}_\psi$ is constructed according to Definition 4.18:

$$\begin{aligned}
\exists_F p.\mathbb{A}_\psi \text{ accepts } \mathbb{T} &\text{ iff } \mathbb{A}_\psi \text{ accepts } \mathbb{T}[p \mapsto X], \text{ for some } X \subseteq_\omega T && \text{(Lemma 4.20)} \\
&\text{ iff } \mathbb{T}[p \mapsto X] \models \psi, \text{ for some } X \subseteq_\omega T && \text{(induction hyp.)} \\
&\text{ iff } \mathbb{T} \models \exists p.\psi && \text{(semantics WMSO)}
\end{aligned}$$

QED

5 From WMSO-automata to WMSO

In what follows, we verify that WMSO-automata capture exactly the expressive power of WMSO on the class of tree models. Since we already proved the direction from formulas into automata (Theorem 4.5), we just have to verify that there is a sound translation going in the other direction. For this purpose, we first introduce a fixpoint extension of first-order logic.

5.1 Fixpoint extension of first-order logic

Let our first-order signature be composed of a set P of monadic predicates (denoted with capital latin letters) and an unique binary predicate R . Analogously to the modal μ -calculus, the fixpoint extension of $\text{FOE}^\infty(P)$ is defined by adding a fixpoint construction clause.

Definition 5.1 The fixed point logic $\mu\text{FOE}^\infty(P)$ is given by:

$$\varphi ::= q(x) \mid R(x, y) \mid x \approx y \mid \exists x.\varphi \mid \exists^\infty x.\varphi \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mu p.\varphi(p, x)$$

where $p, q \in P$, $x, y \in \text{iVar}$; moreover p occurs only positively in $\varphi(p, x)$ and x is the only free variable in $\varphi(p, x)$.

The semantics of the fixpoint formula $\mu p.\varphi(p, x)$ is the expected one. Given a model \mathbb{T} and $s \in T$, $\mathbb{T} \models \mu p.\varphi(p, s)$ iff s is in the least fixpoint of the operator $F_\varphi : \wp(T) \rightarrow \wp(T)$ defined as $F_\varphi(S) := \{t \in T \mid \mathbb{T}[p \mapsto S] \models \varphi(p, t)\}$.

Formulas of μFOE^∞ may be also classified according to their alternation depth as it happens for the modal μ -calculus. The alternation-free fragment of μFOE^∞ is thence defined as the collection of μFOE^∞ -formulas φ without nesting of greatest and least fixpoint operators, i.e. such that, for any two subformulas $\mu p.\psi_1(p, y)$ and $\nu q.\psi_2(q, z)$, predicates p and q do not occur free respectively in $\psi_2(q, z)$ and $\psi_1(p, y)$.

Definition 5.2 Given $p \in \mathbf{P}$, we say that $\varphi \in \mu\text{FOE}^\infty(\mathbf{P})$ is

- *monotone in the predicate p* iff for every LTS \mathbb{T} and assignment g ,

$$\text{if } \mathbb{T}, g \models \varphi \text{ and } \sigma^b(p) \subseteq E, \text{ then } \mathbb{T}[p \mapsto E], g \models \varphi$$

- *continuous in the predicate p* iff for every LTS \mathbb{T} and assignment g there exists some finite $S \subseteq_\omega \sigma^b(p)$ such that

$$\mathbb{T}, g \models \varphi \quad \text{iff} \quad \mathbb{T}[p \mapsto S], g \models \varphi.$$

In the next definition, we provide a definition of the continuous fragment of μFOE^∞ , reminiscent of the one defined in Theorem 3.31.

Definition 5.3 Let $\mathbf{Q} \subseteq \mathbf{P}$ be a set of monadic predicates. The fragment $\mu\text{FOE}^\infty \mathbf{C}_{\mathbf{Q}}(\mathbf{P})$ is defined by the following rules:

$$\varphi ::= \psi \mid q(x) \mid \exists x.\varphi(x) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{W}x.(\varphi, \psi) \mid \mu p.\varphi'(p, x)$$

where $q \in \mathbf{Q}$, $\psi \in \mu\text{FOE}^\infty(\mathbf{P} \setminus \mathbf{Q})$, $p \in \mathbf{P} \setminus \mathbf{Q}$, $\mathbf{W}x.(\varphi, \psi) := \forall x.(\varphi(x) \vee \psi(x)) \wedge \forall^\infty x.\psi(x)$ and $\varphi'(p, x)$ is a formula with only x free such that $\varphi'(p, x) \in \mu\text{FOE}^\infty \mathbf{C}_{\mathbf{Q} \cup \{p\}}(\mathbf{P})$.

Lemma 5.4 If $\varphi \in \mu\text{FOE}^\infty \mathbf{C}_{\mathbf{Q}}(\mathbf{P})$ then φ is continuous in (each predicate from) \mathbf{Q} .

Proof. First, notice that If $\varphi \in \mu\text{FOE}^\infty \mathbf{C}_{\mathbf{Q}}(\mathbf{P})$ then φ is monotone in (each predicate from) \mathbf{Q} . The proof goes then by induction on the complexity of φ . For the all the cases except the fixpoint one, the proof is the same as the one for Lemma 3.32. For $\varphi = \mu p.\varphi'(p, x)$, with $\varphi'(p, x) \in \mu\text{FOE}^\infty \mathbf{C}_{\mathbf{Q} \cup \{p\}}(\mathbf{P})$, the argument is the same as in [6, Lemma 1]. QED

As for the modal μ -calculus, we define the fragment $\mu_c\text{FOE}^\infty$ of μFOE^∞ as the one where the use of the least fixed point operator is restricted to the continuous fragment.

Definition 5.5 The fragment $\mu_c\text{FOE}^\infty(\mathbf{P})$ of $\mu\text{FOE}^\infty(\mathbf{P})$ is given by the following restriction of the fixpoint operator to the continuous fragment:

$$\varphi ::= q(x) \mid R(x, y) \mid x \approx y \mid \exists x.\varphi \mid \exists^\infty x.\varphi \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mu p.\varphi'(p, x)$$

where $p, q \in \mathbf{P}$, $x, y \in \text{iVar}$; and $\varphi'(p, x) \in \mu\text{FOE}^\infty \mathbf{C}_{\{p\}}(\mathbf{P})$ is such that p occurs only positively in φ' and x is the only free variable in φ' .

We now recall a useful property of fixpoint and continuity. Let $\varphi(p, x)$ a formula with only x free. Given a LTS \mathbb{T} , for every ordinal α , we define by induction the following sets:

- $\varphi^0(\emptyset) := \emptyset$,
- $\varphi^{\alpha+1}(\emptyset) := \{s \in T \mid \mathbb{T}[p \mapsto \varphi^\alpha(\emptyset)] \models \varphi(p, s)\}$,
- $\varphi^\lambda(\emptyset) := \bigcup_{\alpha < \lambda} \varphi^\alpha(\emptyset)$, with λ limit.

If φ is monotone in p , it is possible to show that $\varphi^{\beta+1}(\emptyset) = \varphi^\beta(\emptyset)$, for some ordinal β . Moreover, the set $\varphi^\beta(\emptyset)$ is the least fixpoint of F_φ (cf. for instance [1]).

A formula $\varphi(p, x)$ is said to be *constructive* in p if its least fixpoint is reached in at most ω steps, i.e., if for every model \mathbb{T} , the least fixpoint of F_φ equals to $\bigcup_{\alpha < \omega} \varphi^\alpha(\emptyset)$. From a local perspective, this means that a formula $\varphi(p, x)$ constructive in p if for every model \mathbb{T} , every node $s \in T$, whenever $\mu p. \varphi(p, x)$ is true at s , then s belongs to some finite approximant $\varphi^{i+1}(\emptyset)$ of the least fixpoint of F_φ . The next proposition is easily verified:

Proposition 5.6 *Let $\varphi(p, x)$ be a μFOE^∞ -formula with only x free. If $\varphi(p, x)$ is continuous in p , then for every LTS \mathbb{T} , and every node $s \in T$, there is $i < \omega$ such that*

$$\mathbb{T} \models \mu p. \varphi(p, s) \text{ iff } s \in \varphi^{i+1}(\emptyset).$$

From the fact that sets $\varphi^{i+1}(\emptyset)$ are essentially defined as finite unfoldings and the previous Proposition 5.6, we obtain the following.

Proposition 5.7 *Let $\varphi(p, x)$ be a μFOE^∞ -formula with only x free and such that $\varphi(p, x)$ is continuous in p . Let \mathbb{T} be a LTS, and $s \in T$. Then $\mathbb{T} \models \mu p. \varphi(p, s)$ iff there is a finite set $p^\mathbb{T} \subseteq_\omega T$ such that $s \in p^\mathbb{T}$ and $\mathbb{T}[p \mapsto p^\mathbb{T}] \models \varphi(p, t)$ for every $t \in p^\mathbb{T}$.*

Proof. For the direction from left to right, assume that $\mathbb{T} \models \mu p. \varphi(p, s)$. By Proposition 5.6, we know that there is $i < \omega$ such that $\mathbb{T}[p \mapsto \varphi^i(\emptyset)] \models \varphi(p, s)$. The set $\varphi^i(\emptyset)$ need not to be finite. However, using this information, we are going to construct a finite tree whose nodes t are labelled by finite sets X_j^m , where m is a node of \mathbb{T} and $j \leq i$, satisfying the following condition:

1. if t is the root, then t is labelled by X_i^s ,
2. if t is labelled by $X_j^m = \{s_1, \dots, s_\ell\}$ and $j > 0$, then t has ℓ children and for every $s_i \in X_j^m$ there is a unique child t' of t labelled by $X_{j-1}^{s_i}$ where m is a node,
3. for every node t of the tree, if t is labelled by X_j^m , then it holds that $X_j^m \subseteq \varphi^j(\emptyset)$.

If we verify that $\mathbb{T}[p \mapsto p^\mathbb{T}] \models \varphi(p, s)$ holds by taking as $p^\mathbb{T}$ the union of all labels of the nodes of the constructed tree, we can conclude for the proof of this direction.

As starting point of the inductive construction, we start by the empty tree. Recall that we know that $\mathbb{T}[p \mapsto \varphi^i(\emptyset)] \models \varphi(p, s)$. Since $\varphi(p, x)$ is continuous in p , there is a finite set $X_i^s \subseteq \varphi^i(\emptyset)$ such that $\mathbb{T}[p \mapsto X_i^s] \models \varphi(p, s)$. We then add a root to our tree and label it by X_i^s . Assume that at a leaf s of our tree is labelled by X_j^m , for some $j < i$. If X_j^m is

empty, then we stop, else we proceed as follows. We know that $X_j^m \subseteq \varphi^j(\emptyset)$. This means that $\mathbb{T}[p \mapsto \varphi^{j-1}(\emptyset)] \models \varphi(p, r)$, for every $r \in X_j^m$. By continuity, for each such r , there is a finite set $X_{j-1}^m \subseteq \varphi^{j-1}(\emptyset)$ such that $\mathbb{T}[p \mapsto X_{j-1}^m(\emptyset)] \models \varphi(p, r)$. For each $r \in X_j^m$ we thus add a child to m and label it with X_{j-1}^r . By definition of $\varphi^{i+1}(\emptyset)$, the tree is finite. Let X be the union of all labels of the constructed tree. X is finite, and by monotonicity of $\varphi(p, x)$ we have that for every $m \in X \cup \{s\}$, $\mathbb{T}[p \mapsto X \cup \{s\}] \models \varphi(p, m)$.

For the other direction, it's enough to notice that the smallest finite set $p^\mathbb{T} \subseteq T$ such that $\mathbb{T}[p \mapsto p^\mathbb{T}] \models \varphi(p, s)$ and $\mathbb{T}[p \mapsto p^\mathbb{T}] \models \varphi(p, m)$ for all $m \in p^\mathbb{T}$ is the least fixpoint F_φ . QED

Proposition 5.7 naturally suggests the following translation $(-)^* : \mu\text{FOE}^\infty(\mathbf{P}) \rightarrow \text{WMSO}(\mathbf{P})$,

- $(p(x))^* = p(x)$,
- $(R(x, y))^* = R(x, y)$
- $(x \approx y)^* = (x \approx y)$
- $(\varphi \wedge \psi)^* = (\varphi)^* \wedge (\psi)^*$,
- $(\neg \varphi)^* = \neg(\varphi)^*$,
- $(\exists x. \varphi)^* = \exists x. (\varphi)^*$,
- $(\exists^\infty x. \varphi)^* = \forall p. \exists x. (\neg p(x) \wedge (\varphi)^*)$,
- $(\mu p. \varphi(p, x))^* = \exists p(p(x) \wedge \forall y(p(y) \rightarrow (\varphi(p, y))^*))$.

The following theorem is then an immediate corollary of Proposition 5.7.

Theorem 5.8 *For every alternation-free formula φ in the $\mu_c\text{FOE}^\infty$, every LTS \mathbb{T} , and assignment g , we have $\mathbb{T}, g \models \varphi$ iff $\mathbb{T}, g \models (\varphi)^*$.*

Proof. The proof goes by induction on the complexity of φ , the only critical step being the least fixpoint operator one. But this follows by applying Proposition 5.7 and the induction hypothesis. QED

5.2 Translating automata into formulas

We are now ready to prove the second main result of the paper.

Theorem 5.9 *There is an effective procedure that given an automaton in $\text{Aut}_{cw}(\text{FOE}_1^\infty)$, returns an equivalent WMSO-formula.*

Proof sketch. The argument is essentially a refinement of the standard proof showing that any automaton in $\text{Aut}(\text{FO}_1)$ can be translated into an equivalent μ -formula $\xi_\mathbb{A}$ (cf. e.g. [22]). The idea is the following. We see a WMSO-automaton as a system of equations expressed in terms of FOE^∞ -formulas: each state corresponds to a monadic predicate variable and the parity of a state corresponds to the least and greatest fixpoint that we seek for the associated variable, etc. One then solves this system of equations via the same inductive procedure used

to obtain the formula of the modal μ -calculus from the system associated with a $Aut(FO_1)$ -automaton (see e.g. [1] for a description of the solution procedure). Because of the (weakness) and (continuity) conditions on the starting WMSO-automaton \mathbb{A} , it is thence possible to verify that the resulting fixpoint formula $\xi_{\mathbb{A}}$ belongs to μ_cFOE^∞ . QED

Remark 5.10 As a corollary of the automata characterization on trees of WMSO, we obtain the equivalence on this class of structures between WMSO and μ_cFOE^∞ . This consequence should be compared to the analogous result obtained by Walukiewicz in [23] for FPL (fixpoint extension of FOE) and MSO on trees. \triangleleft

6 Modal Characterization of WMSO

In this section we are going to prove Theorem 1.2. Our proof of the first item of the theorem crucially involves automata. In the previous section we saw that on trees, WMSO effectively corresponds to the automata class $Aut_{cw}(FOE_1^\infty)$. We will now relate this class to the one of parity automata based on FO_1 and satisfying similar weakness and continuity conditions.

Definition 6.1 A μ_cML -automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ is an automaton $\mathbb{A} \in Aut(FO_1)$ such that for all states $a, b \in A$ with $a \preceq b$ and $b \preceq a$ the following conditions hold:

(weakness) $\Omega(a) = \Omega(b)$,

(continuity) if $\Omega(a)$ is odd (resp. even) then, for each $c \in C$ we have $\Delta(a, c) \in FO_1^+ C_b(A)$ (resp. $\Delta(a, c) \in FO_1^+ \overline{C}_b(A)$).

As the class of such automata coincides with $Aut_{cw}(FO_1)$ we use the same name to denote it.

As the key technical result of our paper, in subsection 6.1 we will provide a construction $(-)^{\bullet} : Aut_{cw}(FOE_1^\infty) \rightarrow Aut_{cw}(FO_1)$, such that for all \mathbb{A} and \mathbb{T} we have

$$\mathbb{A}^{\bullet} \text{ accepts } \mathbb{T} \text{ iff } \mathbb{A} \text{ accepts } \mathbb{T}^\omega, \quad (12)$$

where \mathbb{T}^ω is the ω -unravelling of \mathbb{T} . As we shall see, the map $(-)^{\bullet}$ is completely determined at the one-step level, that is, by some model-theoretic connection between FOE_1^∞ and FO_1 .

The second fact, to be discussed in subsection 6.2, is that for each μ_cML -automaton \mathbb{A} we can effectively construct an equivalent μ_cML -formula $\xi_{\mathbb{A}}$.

On the basis of the above observations we show that those results are enough to prove Theorem 1.2(i) as follows:

Proof of Theorem 1.2. (1) Given a WMSO-formula φ , let $\varphi^{\bullet} := \xi_{\mathbb{A}_{\varphi}^{\bullet}}$. We verify that φ is bisimulation invariant iff φ and φ^{\bullet} are equivalent. The direction from right to left is immediate by the observation that φ^{\bullet} is a formula in μML . The opposite direction follows from the following chain of equivalences:

$$\begin{aligned} \mathbb{T} \models \varphi \text{ iff } \mathbb{T}^\omega \models \varphi & \quad (\varphi \text{ bisimulation invariant}) \\ \text{iff } \mathbb{A}_{\varphi} \text{ accepts } \mathbb{T}^\omega & \quad (\varphi \equiv \mathbb{A}_{\varphi} \text{ on trees}) \\ \text{iff } \mathbb{A}_{\varphi}^{\bullet} \text{ accepts } \mathbb{T} & \quad (12) \\ \text{iff } \mathbb{T} \models \xi_{\mathbb{A}_{\varphi}^{\bullet}} & \quad (\mathbb{A}_{\varphi}^{\bullet} \equiv \xi_{\mathbb{A}_{\varphi}^{\bullet}}) \end{aligned}$$

(2) For the second part of Theorem 1.2, We first define, for every first-order variable x , a translation ST_x from the μ -calculus into the set of μFOE^∞ -formulas with only x free:

- $ST_x(p) = p(x)$
- $ST_x(\varphi \wedge \psi) = ST_x(\varphi) \wedge ST_x(\psi)$,
- $ST_x(\varphi \vee \psi) = ST_x(\varphi) \vee ST_x(\psi)$,
- $ST_x(\neg\varphi) = \neg ST_x(\varphi)$,
- $ST_x(\Diamond\varphi) = \exists y(R(x, y) \wedge ST_y(\varphi))$,
- $ST_x(\mu p.\varphi) = \mu p.ST_x(\varphi)$,

Clearly, every formula of the $\mu_c\text{ML}$ -fragment of the μ -calculus is mapped to a logically equivalent formula of the $\mu_c\text{ML}$ -fragment of μFOE^∞ . Let $(-)_\bullet : \mu_c\text{ML} \rightarrow \text{WMSO}$ defined as the composite $(-)^* \circ ST_x$. By Theorem 5.8 we obtain that $\psi \equiv \psi_\bullet$, for all $\psi \in \mu_c\text{ML}$. QED

6.1 One-step translations

In this subsection we will define a construction that transforms an arbitrary automaton \mathbb{A} in $\text{Aut}_{cw}(\text{FOE}_1^\infty)$ into an automaton \mathbb{A}^\bullet in $\text{Aut}_{cw}(\text{FO}_1)$, such that \mathbb{A} and \mathbb{A}^\bullet are related as in (12). This construction is completely determined by the following translation at the one-step level.

Definition 6.2 Using the fact that by Corollary 3.26, any formula in $\text{FOE}_1^{\infty+}(A)$ is equivalent to a disjunction of formulas of the form $\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$, we define the translation $(-)^{\bullet} : \text{FOE}_1^{\infty+}(A) \rightarrow \text{FO}_1^+(A)$ as follows. We set

$$\left(\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)\right)^{\bullet} := \bigwedge_i \exists x_i. \tau_{T_i}^+(x_i) \wedge \forall x. \bigvee_{S \in \Sigma} \tau_S^+(x)$$

and for $\alpha = \bigvee_i \alpha_i$ we define $\alpha^{\bullet} := \bigvee \alpha_i^{\bullet}$.

The key property of this translation is the following.

Proposition 6.3 *For every one-step model (D, V) and every $\alpha \in \text{FOE}_1^{\infty+}(A)$ we have*

$$(D, V) \models \alpha^{\bullet} \text{ iff } (D \times \omega, V_\pi) \models \alpha, \quad (13)$$

where V_π is the induced valuation given by $V_\pi(a) := \{(d, k) \mid d \in V(a), k \in \omega\}$.

Proof. Clearly it suffices to prove (13) for formulas of the form $\alpha = \nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$.

\Rightarrow Assume $(D, V) \models \alpha^{\bullet}$, we will show that $(D \times \omega, V_\pi) \models \nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$. Let d_i be such that $\tau_{T_i}^+(d_i)$ in (D, V) . It is clear that the (d_i, i) provide *distinct* elements satisfying $\tau_{T_i}^+((d_i, i))$ in $(D \times \omega, V_\pi)$ and therefore the first-order existential part of α is satisfied. With a similar but easier argument it is straightforward that the existential generalized quantifier part of α is also satisfied. For the universal parts of $\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$ it is enough to observe that, because

of the universal part of α^\bullet , every $d \in D$ realizes a positive type in Σ . By construction, the same applies to $(D \times \omega, V_\pi)$, therefore this takes care of both universal quantifiers.

$\boxed{\Leftarrow}$ Assuming that $(D \times \omega, V_\pi) \models \nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$, we will show that $(D, V) \models \alpha^\bullet$. The existential part of α^\bullet is trivial. For the universal part we have to show that every element of D realizes the positive part of a type in Σ . Suppose not, and let $d \in D$ be such that $\neg \tau_S^+(d)$ for all $S \in \Sigma$. Then we have $(D \times \omega, V_\pi) \not\models \tau_S^+((d, k))$ for all k . That is, there are infinitely many elements not realizing the positive part of any type in Σ . Hence we have $(D \times \omega, V_\pi) \not\models \forall^\infty y. \bigvee_{S \in \Sigma} \tau_S^+(y)$. Absurd, because that is part of $\nabla_{\text{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$. QED

As a consequence of Proposition 6.3 we obtain the following.

Definition 6.4 Given an automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $\text{Aut}(\text{FOE}_1^\infty)$, define the automaton $\mathbb{A}^\bullet := \langle A, \Delta^\bullet, \Omega, a_I \rangle$ in $\text{Aut}(\text{FO}_1)$ by putting, for each $(a, c) \in A \times C$:

$$\Delta^\bullet(a, c) := (\Delta(a, c))^\bullet.$$

Proposition 6.5 For any automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $\text{Aut}(\text{FOE}_1^\infty)$, and any model \mathbb{T} , \mathbb{A} and \mathbb{T} satisfy (12).

Proof. The proof of this proposition is based on a fairly routine comparison of the acceptance games $\mathcal{A}(\mathbb{A}^\bullet, \mathbb{T})$ and $\mathcal{A}(\mathbb{A}, \mathbb{T}^\omega)$. In a slightly more general setting, the details of this proof can be found in [21]. QED

It remains to be checked that the construction $(-)^^\bullet$, which has been defined for arbitrary automata in $\text{Aut}(\text{FOE}_1^\infty)$, transforms WMSO-automata into automata in the right class, viz., $\text{Aut}_{cw}(\text{FO}_1)$.

Proposition 6.6 Let $\mathbb{A} \in \text{Aut}(\text{FOE}_1^\infty)$. If $\mathbb{A} \in \text{Aut}_{cw}(\text{FOE}_1^\infty)$, then $\mathbb{A}^\bullet \in \text{Aut}_{cw}(\text{FO}_1)$.

Proof. This proposition can be verified by a straightforward inspection, at the one-step level, that if a formula $\alpha \in \text{FOE}_1^{\infty+}(A)$ belongs to the fragment $\text{FOE}_1^{\infty+}C_a(A)$, then its translation α^\bullet lands in the fragment $\text{FO}_1^+C_a(A)$. QED

Remark 6.7 As a corollary of the previous two propositions we find that

- $\text{Aut}(\text{FO}_1) \equiv \text{Aut}(\text{FO}_1)/\simeq$, and
- $\text{Aut}_{cw}(\text{FO}_1) \equiv \text{Aut}_{cw}(\text{FO}_1)/\simeq$.

In fact, we are dealing here with an instantiation of a more general phenomenon that is essentially coalgebraic in nature. In [21] it is proved that if \mathcal{L}_1 and \mathcal{L}'_1 are two one-step languages that are connected by a translation $(-)^^\bullet : \mathcal{L}'_1 \rightarrow \mathcal{L}_1$ satisfying a condition similar to (13), then we find that $\text{Aut}(\mathcal{L}_1)$ corresponds to the bisimulation-invariant fragment of $\text{Aut}(\mathcal{L}'_1)$: $\text{Aut}(\mathcal{L}_1) \equiv \text{Aut}(\mathcal{L}'_1)/\simeq$. This subsection can be generalized to prove similar results relating $\text{Aut}_w(\mathcal{L}_1)$ to $\text{Aut}_w(\mathcal{L}'_1)$, and $\text{Aut}_{cw}(\mathcal{L}_1)$ to $\text{Aut}_{cw}(\mathcal{L}'_1)$. \triangleleft

6.2 From automata in $Aut_{cw}(\text{FO}_1)$ to formulas in $\mu_c\text{ML}$

In this subsection we focus on the following theorem.

Theorem 6.8 *There is an effective procedure that, given an automaton \mathbb{A} in $Aut_{cw}(\text{FO}_1)$, returns an equivalent formula $\xi_{\mathbb{A}}$ of the fragment $\mu_c\text{ML}$ of the modal μ -calculus.*

Proof. The argument is a refinement of the standard proof showing that any automaton \mathbb{A} from $Aut(\text{FO}_1)$ can be translated into an equivalent μ -formula $\xi_{\mathbb{A}}$ (cf. e.g. [22]), and it is essentially a special case of the argument proving Theorem 5.9.

From now on, we always assume that a formula $\Delta(a, c)$ is in normal form. Following [22], we introduce another type of automata, called (P, X) -automata, which operate on $\wp(P \cup X)$ -trees. They differ from automata whose one-step language is defined over predicates in $(A \cup X)$ in that⁴

- Monadic predicate letters from X can occur in the scope of a quantifier and only there, meaning that (P, X) -automata have transition $\Delta(a, c) \in \text{FO}_1^+(A \cup X)$
- The transition function is uniquely determined by the restriction of the coloring to P , that is, for every $a \in A$, and $c_1, c_2 \in C$, if $c_1 \cap P = c_2 \cap P$ then $\Delta(a, c_1) = \Delta(a, c_2)$.

We also assume that for every $x \in X$ there is a unique $a \in A$ and an unique $c \in C$ such that x occurs in $\Delta(a, c)$. The notion of acceptance is defined as expected, the only difference with being that during the acceptance game Éloise has to provide a valuation only for predicates in A making formula given by the transition function true. It is then enough to prove the following claim.

Claim 1 *There is an effective procedure that, given a (P, X) -automaton \mathbb{A} gives an equivalent formula $\xi_{\mathbb{A}} \in \mu_c\text{ML}$ in which all occurrences of variables in X are positive.*

PROOF OF CLAIM. Without loss of generality, we can assume that:

- Every (maximal) strongly connected component (SCC) in the graph of \preceq has an unique entrance point,
- The directed acyclic graph (DAG) of the SCCs of \preceq is a tree, and more specifically,
- $\{c \in A \mid a \rightsquigarrow c, c \prec a\} \cap \{c \in A \mid b \rightsquigarrow c, c \prec b\} = \emptyset$ whenever a, b are in the same SCC, with $a \neq b$.

Given a (P, X) -automaton \mathbb{A} , we are now going to define a function $\delta_{\mathbb{A}} : A \rightarrow \text{ML}(A \cup X \cup P)$ that assigns to each state a of \mathbb{A} a modal formula $\delta_{\mathbb{A}}(a)$ over $A \cup X \cup P$ representing all possible transitions from a in the modal language with the property that if $b \in A$ is in the same \preceq -cycle of a and $\Omega(a) = 1$, then $\delta_{\mathbb{A}}(a)$ is continuous in b . Dually for $\Omega(a) = 0$.

Let $c \in C$, and assume $\Delta(a, c)$ is in positive basic form $\bigvee \nabla_{\text{FO}}^+(\Sigma)$. We define a first translation TR_1 taking as argument $\Delta(a, c)$ and giving as result a formula from the (guarded fragment of) first-order logic over $A \cup X \cup P$ as follows. With every disjunct

$$\nabla_{\text{FO}}^+(\Sigma) = \bigwedge_{S \in \Sigma} \exists x. \tau_S^+(x) \wedge \forall z. (\bigvee_{S \in \Sigma} \tau_S^+(z)),$$

⁴Parity automata based on FO_1 are thus simply (P, \emptyset) -automata.

we associate the formula

$$TR_1(\nabla_{FO}^+(\Sigma)) := \bigwedge_{S \in \Sigma} \exists y. (R(x, y) \wedge \tau_S^+(y)) \wedge \forall z. (R(x, z) \rightarrow \bigvee_{S \in \Sigma} \tau_S^+(z)).$$

The formula $TR_1(\nabla_{FO}^+(\Sigma))$ is bisimulation invariant, and it is equivalent to the modal formula

$$TR_2(\nabla_{FO}^+(\Sigma)) := \bigwedge_{S \in \Sigma} \Diamond(\bigwedge S) \wedge \Box \bigvee_{S \in \Sigma} (\bigwedge S).$$

Let $TR_3(\Delta(a, c)) = \bigwedge (c \cap P) \wedge \bigwedge_{p \in P \setminus c} \neg p \wedge \bigvee TR_2(\nabla_{FO}^+(\Sigma))$. The modal formula $\delta_{\mathbb{A}}(a)$ is then defined as

$$\bigvee_{c \in C} TR_3(\Delta(a, c)).$$

By construction we have, for every \mathbb{T} ,

$$\mathbb{T}[x \mapsto s_I] \models \bigvee_{c \in C} (\tau_{(c \cap P)}(y) \wedge \bigvee TR_1(\nabla_{FO}^+(\Sigma))) \quad \text{iff} \quad \mathbb{T} \models \delta_{\mathbb{A}}(a).$$

A modal automaton over P is an automaton $\langle A, \delta, \Omega, a_I \rangle$ such that $\delta : A \rightarrow \text{ML}^+(A)$, where $\text{ML}^+(A)$ is the set of all modal formulas over propositions $A \cup P$ such that elements from A appear only positively. The acceptance game associated with such an automaton and a tree \mathbb{T} is determined by the (symmetric) acceptance game defined according to the rules of Table 2. This means that we can equivalently see the automaton \mathbb{A} as a modal automaton $\langle A, \delta_{\mathbb{A}}, \Omega, a_I \rangle$ whose transition function satisfies the weakness and continuity conditions. Thus, from now on we see \mathbb{A} as the equivalent modal automaton we have just described.

Position	Player	Admissible moves	Parity
$(a, s) \in A \times S$	\exists	$\{(\delta_{\mathbb{A}}(a), s)\}$	$\Omega(a)$
$(\psi_1 \vee \psi_2, s)$	\exists	$\{(\psi_1, s), (\psi_2, s)\}$	—
$(\psi_1 \wedge \psi_2, s)$	\forall	$\{(\psi_1, s), (\psi_2, s)\}$	—
$(\Diamond \varphi, s)$	\exists	$\{(\varphi, t) \mid t \in R[s]\}$	—
$(\Box \varphi, s)$	\forall	$\{(\varphi, t) \mid t \in R[s]\}$	—
$(\neg p, s) \in P \times S$ and $p \notin \sigma(s)$	\forall	\emptyset	—
$(\neg p, s) \in P \times S$ and $p \in \sigma(s)$	\exists	\emptyset	—
$(p, s) \in P \times S$ and $p \in \sigma(s)$	\forall	\emptyset	—
$(p, s) \in P \times S$ and $p \notin \sigma(s)$	\exists	\emptyset	—

Table 2: (Symmetric) acceptance game for modal automata

For the construction of $\xi_{\mathbb{A}}$, we proceed by induction on the tree height of the DAG t of SCC. If the height is 1, that is the DAG is a single point graph, we reason as follows. We have two cases to consider: either the SCC is trivial (i.e. it consists of a single non looping node), or not. In the first case, $A = \{a_I\}$ and \mathbb{A} is equivalent to $\xi_{\mathbb{A}} := \delta_{\mathbb{A}}(a_I)$.

For the second case, let us assume that $\Omega(a_I) = 1$, the case when it is 0 being, mutatis mutandis, the same. Let $A = \{a_0, \dots, a_\ell\}$, and $a_I = a_\ell$. Since \mathbb{A} is a weak modal automaton

satisfying the continuity condition, given $a, b \in A$, if b occurs in $\delta_{\mathbb{A}}(a)$, then b is only in the scope of \Diamond operator. We can now see the automaton \mathbb{A} as a system of modal equations, and by applying the standard inductive procedure *solves* this system of equations and construct the least fixpoint formula $\xi_{\mathbb{A}}$ equivalent to \mathbb{A} ,

Here the key observation is that the weakness and continuity conditions on strongly connected components of the automaton ensure that when we execute a single step in solving the system of equations, we may work within the (syntactically) continuous fragment of the modal μ -calculus. From this, we can deduce that $\xi_{\mathbb{A}} \in \mu_c\text{ML}$. Clearly the procedure preserves the polarity of each $x \in X$, meaning that all variables in X are positive in $\xi_{\mathbb{A}}$.

For the induction step, assume the successors of the root of t are (n_1, \dots, n_ℓ) . For each $i \in \{1, \dots, \ell\}$, let a_i be the entrance point of the SCC of n_i , and let \mathbb{A}_i be the automaton \mathbb{A} but having as an initial state a_i . If we do not consider the states that are not reachable by a_i , the DAG of the SCC of \mathbb{A}_i is $t.n_i$ (the subtree of t starting at n_i). Let $Y = \{a_1, \dots, a_\ell\}$, and M be the set of states of \mathbb{A} that belong to the root of t . We assume $X \cap Y = \emptyset$. The structure

$$\mathbb{A}_M := \langle M, \delta_{\mathbb{A}}|_M, \Omega|_M, a_I \rangle$$

is a $(P, X \cup Y)$ -automaton.

The inductive hypothesis applies to automata $\mathbb{A}_M, \mathbb{A}_1, \dots, \mathbb{A}_\ell$. Thus we obtain fixpoint formulas $\xi_M, \xi_1, \dots, \xi_\ell$, the former taking free variables in $P \cup X \cup Y$, all the remaining in $P \cup X$, equivalent to $\mathbb{A}_M, \mathbb{A}_1, \dots, \mathbb{A}_\ell$ respectively. Notice that:

- by induction hypothesis, $\xi_M, \xi_1, \dots, \xi_\ell \in \mu_c\text{ML}$, every variable $x \in X \cup Y$ is positive in each of those formulas (if x occurs in it),
- by construction, if x is free in ξ_i , then x is not bounded in ξ_M .

We can therefore deduce that $\xi_{\mathbb{A}} = \xi_M[a_1 \mapsto \xi_1, \dots, a_\ell \mapsto \xi_\ell] \in \mu_c\text{ML}$, and that each variable from X occurs positively in $\xi_{\mathbb{A}}$.

We verify that $\mathbb{T} \models \xi_M[a_1 \mapsto \xi_1, \dots, a_\ell \mapsto \xi_\ell]$ iff $\mathbb{T} \in \mathcal{T}(\mathbb{A})$, for every model \mathbb{T} . But this follows by the following two facts:

1. $\mathbb{T} \models \xi_M[a_1 \mapsto \xi_1, \dots, a_\ell \mapsto \xi_\ell]$ iff $\mathbb{T}[a_1 \mapsto \llbracket \xi_1 \rrbracket^{\mathbb{T}}, \dots, a_\ell \mapsto \llbracket \xi_\ell \rrbracket^{\mathbb{T}}] \models \xi_M$,
2. $\mathbb{T} \in \mathcal{T}(\mathbb{A})$ iff $\mathbb{T}[a_1 \mapsto \llbracket \mathbb{A}_1 \rrbracket^{\mathbb{T}}, \dots, a_\ell \mapsto \llbracket \mathbb{A}_\ell \rrbracket^{\mathbb{T}}] \in \mathcal{T}(\mathbb{A}_M)$

where $\llbracket \mathbb{A}_i \rrbracket^{\mathbb{T}} := \{s \in \mathbb{T} \mid \mathbb{T}.s \in \mathcal{T}(\mathbb{A}_i)\}$. ◀

This finishes the proof of Theorem. QED

7 Conclusions

Overview. In this work we have presented three main contributions. First, we proved that the bisimulation-invariant fragment of WMSO is the fragment $\mu_c\text{ML}$ of μML where the application of the fixpoint operator μp is restricted to formulas that are continuous in p . Our result sheds light on the relationship between MSO, WMSO and μML . In particular, it

provides a positive answer to the question whether $\text{WMSO}/\equiv \leq \text{AFMC}$ on trees of arbitrary branching degree, left open in [5]. This may also be read as the statement that the formulas that separate WMSO from MSO are not bisimulation invariant (and hence, irrelevant in the sense mentioned in the introduction).

To achieve this result, we shaped WMSO-automata, a special kind of parity automata satisfying additional *continuity* and *weakness* conditions, with transition map given by the monadic logic FOE_1^∞ . Our second main contribution was to show that they characterize WMSO on tree models.

As our third main contribution we gave a detailed model-theoretic analysis of the monotone and continuous fragments of FOE_1^∞ . We provide strong normal forms and syntactic characterizations that, besides being of independent interest, are critical for the development of the aforementioned results.

Future work. A first line of research is directly inspired by the methods employed in this work. WMSO-automata and $\mu_c\text{ML}$ -automata are essentially obtained by imposing conditions on the appropriate one-step logic \mathcal{L}_1 and transition structure of automata belonging to $\text{Aut}(\mathcal{L}_1)$. Following this approach, one could take aim at the automata-theoretic counterpart of other fragments of the modal μ -calculus, like PDL, CTL or CTL*.

Another direction of investigation is based on the observation that, from a topological point of view, all WMSO-definable properties are Borel. Since we do not have examples of Borel MSO-definable properties that are not WMSO-definable, it is tempting to conjecture that WMSO is the Borel fragment of MSO and analogously for $\mu_c\text{ML}$ and μML .

Acknowledgements. The second author is supported by the *Expressiveness of Modal Fixpoint Logics* project realized within the 5/2012 Homing Plus programme of the Foundation for Polish Science, co-financed by the European Union from the Regional Development Fund within the Operational Programme Innovative Economy (“Grants for Innovation”). The fourth author is supported by the project ANR 12IS02001 PACE.

References

- [1] André Arnold and Damian Niwiński. *Rudiments of μ -calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, 2001.
- [2] Anuj Dawar and Martin Otto. Modal characterisation theorems over special classes of frames. *Ann. Pure Appl. Logic*, 161(1):1–42, 2009.
- [3] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Perspectives in Mathematical Logic. Springer, 1995.
- [4] Ernest Allen Emerson and Charanjit S. Jutla. Tree automata, μ -calculus and determinacy (extended abstract). In *FOCS*, pages 368–377. IEEE Computer Society, 1991.

- [5] Alessandro Facchini, Yde Venema, and Fabio Zanasi. A characterization theorem for the alternation-free fragment of the modal μ -calculus. In *LICS*, pages 478–487. IEEE Computer Society, 2013.
- [6] Gaëlle Fontaine. Continuous fragment of the mu-calculus. In Michael Kaminski and Simone Martini, editors, *CSL*, volume 5213 of *Lecture Notes in Computer Science*, pages 139–153. Springer, 2008.
- [7] Gaëlle Fontaine and Yde Venema. Some model theory for the modal μ -calculus: syntactic characterizations of semantic properties. Submitted, 2012.
- [8] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
- [9] David Janin and Igor Walukiewicz. On the expressive completeness of the propositional μ -calculus with respect to monadic second order logic. In *Proceedings of the 7th International Conference on Concurrency Theory, CONCUR '96*, pages 263–277, London, UK, 1996. Springer-Verlag.
- [10] Christian Kissig and Yde Venema. Complementation of coalgebra automata. In Alexander Kurz, Marina Lenisa, and Andrzej Tarlecki, editors, *CALCO*, volume 5728 of *Lecture Notes in Computer Science*, pages 81–96. Springer, 2009.
- [11] Phokion G. Kolaitis and Jouko A. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic*, 74(1):23 – 75, 1995.
- [12] Andrzej Mostowski. On a generalization of quantifiers. *Fundamenta Mathematicae*, 44(1):12–36, 1957.
- [13] Andrzej Mostowski. Games with forbidden positions. Technical Report 78, University of Gdansk, 1991.
- [14] David E. Muller, Ahmed Saoudi, and Paul E. Schupp. Alternating automata, the weak monadic theory of trees and its complexity. *Theor. Comput. Sci.*, 97(2):233–244, 1992.
- [15] Damian Niwinski. On fixed-point clones (extended abstract). In Laurent Kott, editor, *ICALP*, volume 226 of *Lecture Notes in Computer Science*, pages 464–473. Springer, 1986.
- [16] David Park. On the semantics of fair parallelism. In *Proceedings of the Abstract Software Specifications, 1979 Copenhagen Winter School*, pages 504–526, London, UK, UK, 1980. Springer-Verlag.
- [17] Eric Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.

- [18] Balder ten Cate and Alessandro Facchini. Characterizing EF over infinite trees and modal logic on transitive graphs. In Filip Murlak and Piotr Sankowski, editors, *MFCS*, volume 6907 of *Lecture Notes in Computer Science*, pages 290–302. Springer, 2011.
- [19] Jouko Väänänen. Remarks on generalized quantifiers and second-order logics. In *Set theory and hierarchy theory*, pages 117–123. Prace Naukowe Instytutu Matematyki Politechniki Wrocławskiej, Wrocław, 1977.
- [20] Johan van Benthem. *Modal Correspondence Theory*. PhD thesis, Universiteit van Amsterdam, 1977.
- [21] Yde Venema. Expressiveness modulo bisimilarity: a coalgebraic perspective. To appear in: A. Baltag and S. Smets (eds.), *Johan van Benthem, Outstanding Contributions Series*, Springer., 2014.
- [22] Yde Venema. Lectures on the modal μ -calculus. Lecture Notes, Institute for Logic, Language and Computation, University of Amsterdam, October 2012.
- [23] Igor Walukiewicz. Monadic second order logic on tree-like structures. In Claude Puech and Rüdiger Reischuk, editors, *STACS*, volume 1046 of *Lecture Notes in Computer Science*, pages 401–413. Springer, 1996.
- [24] Fabio Zanasi. Expressiveness of monadic second order logics on infinite trees of arbitrary branching degree. Master’s thesis, ILLC, Universiteit van Amsterdam, the Netherlands, 2012.